AD-A154 139 ADAPTIVE SIGNAL DETECTION AND OTHER TOPICS IN REAL TIME SIGNAL PROCESSING(U) NAVAL SURFACE WEAPONS CENTER SILVER SPRING MD R H BARAN SEP 84 NSWC/TR-84-412 1/1

UNCLASSIFIED F/G 17/2 NL

END

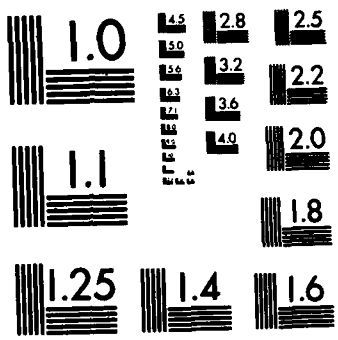

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

NSWC TR 84-412

ADAPTIVE SIGNAL DETECTION AND OTHER TOPICS IN REAL TIME SIGNAL PROCESSING

BY R. H. BARAN

UNDERWATER SYSTEMS DEPARTMENT

SEPTEMBER 1984

Approved for public release, distribution unlimited.

NAVAL SURFACE WEAPONS CENTER

Dahlgren, Virginia 22448 • Silver Spring, Maryland 20910

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER NSWC TR 84-412	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) ADAPTIVE SIGNAL DETECTION AND OTHER TOPICS IN REAL TIME DATA PROCESSING		5. TYPE OF REPORT & PERIOD COVERED
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Robert H. Baran		8. CONTRACT OR GRANT NUMBER(s)
9. PERFORMING ORGANIZATION NAME AND ADDRESS Naval Surface Weapons Center (Code U12) White Oak Silver Spring, MD 20910		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 61152N; ZR00001; U01AA
11. CONTROLLING OFFICE NAME AND ADDRESS		12. REPORT DATE September 1984
		13. NUMBER OF PAGES 79
14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office)		15. SECURITY CLASS. (of this report) UNCLASSIFIED
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release, distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Signal Processing
Data Processing
Communications
Radio
Statistics

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

—Communications receivers are commonly optimized with respect to a stationary interference source. When the interference originates in a composite source, the optimal receiver may be represented as a collection of generically similar subreceivers each one of which is optimized with respect to one of the interference subsources. Selection of the correct subreceiver is conditioned on identifying the subsource that is interfering at a given time. This selection process is controlled by a statistical -

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

S/N 0102-LF-014-6601

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

20. (Cont.)

demultiplexer. In the absence of real communications, when the statistical demultiplexer operates on uncorrelated noise, its function is to estimate the present state of a spectrally white, non-Gaussian, nonstationary random process. To evaluate the performance of the statistical demultiplexer, it is necessary to assert that the state of the observed process is itself a stationary (unobservable) process. Consideration is given to the case in which the unobservable is Markov.

When the receiver is an isolated system that draws power from a limited reservoir of energy, and messages arrive only sporadically, the statistical demultiplexer in its simplest form detects the mere presence of the message and controls a switch in the power line to the main receiver subsystem. The alertness strategy embodied in this scheme can greatly extend the lifetime of the system under conditions that are considered in some detail.

FOREWORD

The Naval Surface Weapons Center's Independent Research (IR) Program is designed to stimulate original work of a basic nature and increase competence in all fields of science and technology relevant to the Center's mission.

This report documents several facets of the engineering studies performed by the author under an IR project titled, "Adaptive Signal Detection," from January 1982 through September 1984.

Approved by:

G. P. KALAF, Head
Underwater Weapons Division

CONTENTS

ILLUSTRATIONS

TABLES

CHAPTER 1

ADAPTIVE SIGNAL DETECTION

The problem of detecting a broadband AM communication signal in the additive noise channel has an optimal solution dictated by the Neyman-Pearson lemma. The signal has the sampled data representation

$$s(1),\ s(2),\ \ldots,\ s(N)$$

so that the input (voltage) at the n-th instant is

$$X(n) = s(n) + Y(n)$$

when Y is the interference contribution. Synchronicity is assumed in order to keep the problem statement tractable. If the expected value of the product $s(n)s(m)$ is zero except when n is m, the signal is said to be uncorrelated. Assume that both the signal and the noise are nearly uncorrelated. Then the only missing ingredient in the formulation of the receiver design problem is the probability density function (pdf) of Y. If $p(x)$ is the density in question, then

$$L\{X\} = \sum_{n=1}^{N} \{\log p[X(n) - s(n)] - \log p[X(n)]\}$$

is the logarithm of the likelihood ratio statistic prescribed by the lemma; and the receiver asserts that the signal is present when L exceeds a threshold. The last summation results from factorization of the likelihood ratio itself. This factorization, in turn, implies that the sequence of Y's is a sequence of independent, identically distributed (iid) random variables. Mathematical

statisticians will point out that lack of correlation is not sufficient to prove independence except when p is the Gaussian density. Thus the solution is not truly rigorous. When Y is in fact Gaussian with zero mean and unit variance, the L-statistic reduces to a sum of the terms

$$X(n)s(n) - (1/2)s^2(n)$$

over all n. The algorithm that computes this sum is the sampled data equivalent of the matched filter first defined by North in the 1940's.

The problem just described may be categorized as the detection of a known signal in spectrally white noise. The matched filter solves the problem when the white noise is also Gaussian. When the noise is non-Gaussian, the L-statistic may be evaluated directly; or, under the assumption that the r.m.s. signal is very small compared to the r.m.s. noise, that statistic can be approximated by the sum of terms

$$g[X(n)]s(n) - (1/2)s^2(n),$$

where

$$g(x) = -d[\log p(x)]/dx,$$

as shown by Antonov.[1] The proof consists in writing the Taylor series for L and simplifying with the aid of several assumptions. The data processing structure implied by the last sum is a no-memory nonlinearity (g) followed by a matched filter.

It may be a little difficult to rationalize the foregoing in light of engineering practice in the field of radio receiver design. Indeed, if the band in question is in the VHF regime, the delay between acquisition of X(n) and X(n+1) should be on the order of several nanoseconds. To support real time operation, one needs a data processor that can convert X to g(X), multiply this number by the appropriate s, and increment the sum before too many nanoseconds elapse. Perhaps some mainframe computers can do this; but the mainstream of communications engineering does not flow through the realm of mainframes. Yet

if the band is in the audio frequency regime or lower, it would seem that almost any microprocessor can handle the job, at least if it is augmented with a coprocessor to perform the multiplication steps.

Evans and Griffiths considered using Antonov's receiver for the reception of ELF signals in submarine communications.[2] Their work was published in 1974, prior to the widespread integration of microcomputers into communications technology. After studying the amplitude probability distribution of ELF atmospheric noise, they showed that the effective enhancement of the signal-to-noise ratio (SNR) attained by placing g(X) in front of the matched filter,

$$e = \int g^2(x)p(x)dx, \tag{1}$$

is typically less than 6 dB.* Moreover, the effective SNR enhancement relative to that achievable by placing a hard limiter there was typically between 0.5 and 3.0 dB. The hard limiter is obtained by setting g(X) equal to the sign of X; and it would coincide with the true g(x) if and only if the noise amplitude has the Laplace distribution. A few dB did not seem like much reward for devising an electronic device exhibiting the peculiar i/o characteristic prescribed by the theory in light of the measured distribution. Also, since the distribution itself is subject to change as the atmospheric conditions evolve from hour to hour, it is not even plausible to suppose that a single device will attain optimality in any meaningful sense.

The optimality theory presumes foreknowledge of a stationary first order density p(x). This knowledge allows the designer to construct a detector with the appropriate $g(x) = -d[\log p(x)]/dx$. When the receiver is a microcomputer, the nonlinear element is just a lookup table in which x is translated into an address. The computer fetches the contents g of the address specified by x and then substitutes g for x in the "matched filter subroutine." Yet a question arises concerning the effect of temporal variation in the density of the noise

---

*When the noise does not have unit variance, the integral is multiplied by the variance to give a dimensionless ratio (e).

amplitude. This question concerns the robustness of likelihood ratio test on which the theory is based, as a statistical test or estimator is said to be robust if it is not too sensitive to flaws in the underlying model. No satisfactory answer of broad generality is readily apparent. Consequently, the designer may address the changing interference environment of the receiver by providing subsystems which allow adaptation. In other words, the optimal nonlinear receiver is inherently an adaptive device which adjusts the shape of the nonlinear element to account for the present properties of the noise.[3]

For the most part, the academic community seems to have rejected adaptive detection, in the sense in which that term is understood here, and turned its attention to so-called nonparametric and distribution-free statistical tests (and the receiver algorithms they engender). The field of nonparametric detection has attracted increasing attention in recent years.[4] Nonparametric detectors are based on statistical hypothesis testing principles for situations where parametric statistical models cannot be specified for the observation under the null hypothesis. (I.e., the density p is unknown, either for lack or data, or because it is nonstationary.) A nonparametric formulation of the problem generally defines a class of allowable distribution functions that cannot be indexed or parameterized. A distribution-free procedure is one based on a statistic, computed from the observations, whose distribution is independent of the precise form of the distribution of the observation.

The neglect accorded to the adaptive procedure seems understandable for two reasons. First, the best nonparametric and distribution-free tests are shown by various authors to be competitive with the corresponding optimal test in terms of asymptotic relative efficiency and other performance measures. For example, in the case of the ELF receiver, only a few more dB of SNR enhancement are gotten by using the locally most powerful (Antonov) test instead of the sign test (the hard limiter) to discern the weak signal. The sign test is an example of a nonparametric procedure.

A second reason for turning away from the adaptive approach lies in the fact that nonstationary random processes are poorly understood. To the communications theoretician, the changing environment can be described as a (sequentially) composite interference source. The optimal receiver is then

represented as a collection of generically similar subreceivers each one of which is optimized with respect to one of the interference subsources. Selection of the correct subreceiver is contingent on identifying the subsource that is interfering at the present time. Let the receiver subsystem that controls the selection process be called the statistical demultiplexer.

CHAPTER 2

STATISTICAL DEMULTIPLEXING

The term "statistical demultiplexing" is derived from the "statistical multiplexer" introduced recently by the Bell system. Within the usual context of time-division multiplexing of signals from several sources through a single channel, the statistical multiplexer checks to see if a given user is talking before routing him to his destination. If a statistical test points to the conclusion that only noise is present, the device skips this user and goes on to consider the next. Now statistical demultiplexing would refer to the case in which the multiplexed data are sent to a receiver without any identifying labels to say which "packet" goes to which destination. The statistical demultiplexer (sdmux) tries to surmise the correct destination on the basis of evidence contained in the given subsequence of the data stream.

The mathematical problem which underlies the design of the sdmux is to block serial data into homogeneous segments. The definition of homogeneity, in the context of the preceding paragraph, involves the question of origin. Such a question might face a hypothetical eavesdropper who has tapped a communications line through which scripted messages are sent through a single type of modem. The common hardware dictates a fixed alphabet to be adopted by all users regardless of content, dialect, or language. If every string of symbols that is bound from user A to user B is headed by a code which (at least) specifies B, then the eavesdropper accomplishes the objective by decoding the headers.

When the composite source is generated not by intelligent users (who have established a neat convention for the routing of messages) but by natural random processes (such as atmospheric interference), the problem reveals its statistical aspect. Now the objective is to establish a microscopic perspective on an observation process in which serial data are blocked according to type and type-casting is by stochastic equivalence. A comprehensive and mathematically

defensible theory of this nature, which seems not to reside in any one place today, is prerequisite to the elaboration of any precise theory of adaptive signal detection in interference environments that are nonstationary in the naive sense.

An extremely simplified model will serve to illustrate the idea. The analyst determines that the noise voltage in a given band, when sampled at a given rate, produces a sequence X(1), ..., X(n), ... of uncorrelated random variables which, for lack of a better hypothesis, are held to be Gaussian with zero mean (although only the mean can be known for sure at the outset). A sample consisting of the first N values of X^2 is used to derive the maximum likelihood estimate of the variance, which is V(1). The same procedure is repeated for the second block of N observations, yielding the variance estimate V(2). The process is continued until last block is ascribed the variance V(L). When L is large, the values of V will exhibit a normal distribution with variance inversely proportional to NI(varX), where varX is the true variance and I is the Fisher information, under the assumption that the observations are homogeneous (that is, they are iid). (Recall that the minimum variance given by the Cramér-Rao lower bound is attained by the complete sufficient statistic which in this case coincides with the maximum likelihood estimator.) But the analyst discovers that V is distributed much more broadly, and that it clusters around two values, v(a) and v(b). Proceeding on a hunch, the analyst lumps all the observations which yielded values of V close to v(a) into a subpopulation $\{X|a\}$ whose complement is $\{X|b\}$. Using one of many well known tests for goodness of fit to a normal distribution with known mean and variance, it is ascertained that the elements of $\{X|a\}$ compromise a sample from a normal population with mean zero and variance v(a) to the satisfaction of some stringent criterion; but the elements of $\{X|b\}$, subjected to the same kind of test, clearly fail to fit the Gaussian model. The analyst now states his conclusion: The noise source has two states or modes which manifest themselves in two kinds of random variables. In the a-state, the sequence of X's is a white Gaussian process with mean power proportional to v(a). In the b-state, the process is non-Gaussian. The analysis engenders a design approach: When the interference source is in the a-state, the optimal receiver has the matched filter structure, since the matched filter is the communications-theoretic counterpart of the likelihood ratio test. When the interference source is in

the b-state, some other kind of test will be more appropriate. A subjective evaluation of many recordings of type b noise finally leads to the conclusion that it is characterized by a Laplace distribution,

$$p(x|b) = \exp(-|x|/c)/2c, \quad c^2 = v(b)/2. \tag{2}$$

But

$$-d\log p(x|b)/dx = (1/c)\operatorname{sgn}(x) \tag{3}$$

is g(x), which defines the shape of the optimal limiter which prefaces the matched filter for signal reception in Laplace noise. Therefore, when the state is b, the optimal receiver performs the sign test.

The sdmux is the receiver subsystem which decides whether the interference source is in state a or state b. Its decision then determines which structure the signal detector assumes. Without the sdmux, the averaged effective SNR at which the receiver operates will be given by

$$P(a)R(a) + P(b)R(b),$$

where P(a) is the fraction of the total time that the interference source is in the a-state, R(a) is the SNR in the a-state, and similiarly for the b-state in the second term. Including the sdmux in the receiver system, the average effective SNR becomes

$$P(a)R(a) + P(b)e(b)R(b),$$

with e(b) = 2 by virtue of Equations (1), (2), and (3). Thus, for example, if the source spends equal time in the two states, and the intensity of the type-b noise is three times the intensity of the type-a noise, the effective average SNR improvement due to the sdmux is

$$\frac{R(a) + 2R(b)}{R(a) + R(b)}$$

or 5/4. Similar arithmetic would show that the adaptive receiver (the one that employs the sdmux) also performs better in the average effective SNR sense than a fixed receiver that performs the sign test (which is optimal only for case b).

The problem with using effective SNR as the performance criterion is that it is not observable independent of the modelling assumptions which underlie the derivation. Real world testing of the adaptive receiver would show up differences in the detection probabilities of the actual symbols which are sent through the channel and lead more naturally to a comparison in terms of the information gain. But to calculate the information gain from basic principles seems to be a much more tedious computational chore.

On a more general and abstract level, the composite noise source is a bivariate process $\{(X,Y)[n],\ n = 1, 2, \ldots\}$ in which each datum $X(n)$ has attached to it a state parameter $Y(n)$ which identifies the distribution from which the datum is drawn. The sequence $\{Y(n),\ n = 1, 2, \ldots)$ is called the <u>side information</u>. The side information accompanies the data when the human-engineered communications source is considered, as noted above; but when nature controls the multiplexing operation (i.e., subsource selection), the side information is <u>not transmitted</u>. Hence, the sdmux has the task of <u>extracting the side information from the data</u> (in order to refer the data to the best subreceiver).

Now the joint density of the data conditioned on the side information is

$$p(\{x\}|\{Y\}) = \prod_{n} p[x(n)|Y(n)] \tag{4}$$

which is the density of a sequence of <u>conditionally independent</u> observations. The information-theoretic problem of deriving the quantity of side information that the data contain presents itself naturally.

$$I(\{X\};\{Y\}) = H(\{X\}) - H(\{X\}|\{Y\})$$

or, since the information is mutual,

$$I(\{X\};\{Y\}) = H(\{Y\}) - H(\{Y\}|\{X\}).$$

Since the sequence {Y} reduces but does not resolve the a priori uncertainty pertaining to {X}, the converse holds true; and the data cannot, on the most fundamental grounds, enable the analyst (or the sdmux) to exactly determine the side information. Because the extraction of side information is inherently imperfect, the performance improvement calculated above for the simple two-state example is properly regarded as an upper bound on what can be achieved using an sdmux which is uninformed about the side information. It also represents the performance achieved when the side information is provided to the receiver's sdmux by some external agent.

CHAPTER 3

THE MEAN ENTROPY RATE OF A MARKOV COMPOSITE SOURCE

In order to clarify the notation of the previous chapter, recall that the entropy of an N-sequence X(1), ..., X(N) of iid observations is

$$NH(X) = -N \sum_{x} p(x) \log p(x) \tag{5}$$

and H(X) is the mean entropy rate (m.e.r.) of a discrete memoryless source that emits data according to the distribution p(x). The m.e.r. of a generalized discrete stationary source is

$$\text{m.e.r.}\{X\} = -\lim (1/N) \sum_{\{x\}} p(\{x\}) \log p(\{x\})$$

when the limit exists as N goes to infinity. The last expression cannot be evaluated directly because the joint density of the N observations indicated by the symbol {X} is given by

$$p(\{x\}) = \sum_{\{y\}} p(\{x\} | \{y\}) p(\{y\})$$

with reference to Equation (4), in which p(.) is the generic probability operator commonly used in elementary expositions of information theory. The m.e.r. is thus the limit of a sum which involves the logarithm of a sum. The usual inequality theorems do not lead to very good bounds on the m.e.r. It may be noted that, if {Y} is an iid sequence, then the last sum can be factored into a product of N terms of the form

$$\sum_{y} p[x(n) | y] p(y) = r[x(n)],$$

with the right hand side being defined by the equality. Then, the observations themselves are iid with distribution $r(x)$; and the entropy of the N-sequence is given by Equation (5) after substituting r for p. The m.e.r. in this case is simply $H(x)$ defined in that manner.

Now the next higher level of complexity must be presented by the case in which $\{Y(n), n = 1, 2, \ldots\}$ is a homogeneous aperiodic Markov chain on a set of discrete states. The formal statement would be

$$P[Y(n) = y \mid Y(n - 1), \ldots, Y(1), X(n - 1), \ldots, X(n)]$$
$$= P[Y(n) = y \mid Y(n - 1)] \ ,$$

wherein the probability that Y assumes the value y at time n is influenced by the history of the bivariate process only through $Y(n - 1)$, its value at the previous instant. Then for a randomly selected n the matrix of elements

$$P[Y(n) = y' \mid Y(n - 1) = y] = t(y' \mid y)$$

is called the transition matrix ($\underline{T}$). When $\underline{T}$ is ascribed an element for every pair (y',y) of states, the Markov chain is fully characterized.

The m.e.r. of the Markov chain itself can be calculated from the transition matrix. Indeed, one has that

$$\text{m.e.r.}\{Y\} = \lim H[Y(N) \mid Y(N - 1)]$$

by analogy with the discrete stationary source of Gallager.[5] This limit is merely

$$\text{m.e.r.}\{Y\} = - \sum_{\{y\}} \sum_{\{y'\}} t(y' \mid y)q(y)\log t(y' \mid y)$$

with $q(y)$ the so-called stationary distribution which is calculated from the transition matrix.

The m.e.r. of the data emitted by the Markov composite source, however, was unspecified by Gallager, who cites (in passing) some investigations by Blackwell in this regard. Since the 1960's, the question seems to have been widely ignored. What would be particularly useful is an expression for m.e.r.{X} subject to restrictions on the transition matrix and its trace. These restrictions would imply that the mean holding time of the Y-process is "long" and that actual transitions are well spaced. Then there are many observations typically intervening between the times that the Y-process enters and exits in a given state.

Suppose that the sequence Y(1), Y(2), ... holds in a given state up to Y(n) which marks the first transition. Now reset the index to one and repeat until the process is renewed again for a different n. Each renewal generates an n(k) and the whole sequence of renewals is n(1), n(2), etc. This sequence is denoted by {n(k)} so that the process holds in the first state from time one to time n(1) - 1, in the second state from n(1) to n(1) + n(2) - 1, etc. At each instant the Markov composite source (MCS) emits an X which corresponds to an integer (since the source is discrete). For a given N-sequence {Y} of states, the N-sequence {X} of observations contains n(j) occurrences of the j-th symbol. Looking only at the subsequence of {X} which occurred in the k+1st phase of {Y}, there were n(j,k) observations of the j-th kind. It can be shown that the total number of possible observation sequences subject to these distributional constraints is

$$W = \prod_{k} [n(k)!/\prod_{j} n(j,k)!] \triangleq \prod_{k,j} w(k,j) .$$

If the probability of having j at a random time within the k-th phase is p(j|k) a priori, then n(j,k) is a multinomial random variable; and the average permutability (or dynamic weight) of N observations is given by

$$EW = \sum_{[n]} \prod_{k,j} w^2(k,j) \exp[n(j,k)\log p(j\ k)]$$

in which the sum covers every matrix of elements n(j,k) consistent with the given numbers {n(k)}.

Now the usual argument made in the statistical analysis of physical ensembles is that, in the limit of large populations, the expected value of the permutability coincides with the modal (or most probable) value. Proofs or derivations of this type of result all flow from the inspiration behind Boltzmann's celebrated H-theorem. If one accepts this idea, the expected value of W is the same as the mode of the distribution of W; and the density of W attains its maximum in the vicinity of those [n] which yield the largest numbers in the summand of the expression for EW. With the usual (Stirling) approximations to the logarithms of the factorial numbers implied in each w(.,.), after some simplification, it becomes apparent that

$$\log W = N \sum_{j,k} Q(j|k)Q(k) \log [p(j|k)/Q^2(j|k)] \tag{6}$$

in which

$$Q(k) = n(k)/n$$

and $Q(j|k)$ is understood as the normalized number of observations of type j in a sample of size n(k) drawn from a population which has the distribution $p(j|k)$. It is essential to note that logW, stated in this way, is subject to a straightforward (though not necessarily easy) computation using analytic principles or Monte Carlo methods. Indeed, if all the n(k)'s tend to infinity, the sample distributions $Q(j|k)$ all approach the corresponding distributions $p(j|k)$; and (logW)/N converges to the average conditional entropy $H(X|Y)$ (after recognizing that a given state Y is revisited an infinite number of times in the same limit).

To conclude the argument, recall that Equation (6) is predicated on a particular realization of the Markov chain. For every particular realization, W attains the same maximum. Entropy is regarded as the logarithm of the permutability. Therefore, the m.e.r. of the MCS must be given by

$$\text{m.e.r.}\{X\} = \lim(1/N)\log W + \text{m.e.r.}\{Y\} \tag{7}$$

in accordance with Equation (6) and the expression above for the mean entropy rate of the Markov process alone. Equation (7) is an approximation that applies when the Y-process is characterized by a long average holding time in every state. Moreover, as the holding times become extremely long, the expression tends to

$$\text{"lim" m.e.r.}\{X\} = H(X|Y) + \text{m.e.r.}\{Y\} \; ;$$

and the approach is from below, since

$$\sum_{j,k} Q(j|k)Q(k)\log[p(j|k)/Q(j|k)] = -K \tag{8}$$

is non-positive. To state this result in words, the mean entropy rate of the data tends toward the sum of the average uncertainty about the observation given the state of the process and the average uncertainty regarding the state given the state at the preceding instant.

CHAPTER 4

VERIFIABILITY AS A RATE-DISTORTION PROBLEM

Equation (7) and the "limit theorem" to which it leads imply that the data convey side information at an average rate

$$\lim\ (1/N)I(\{X\};\{Y\}) = \text{m.e.r.}\{X\} - H(X|Y) = \text{m.e.r.}\{Y\} - K \qquad (9)$$

where K, defined in Equation (8), is the average Kullback information or directed divergence of the sample distribution of a given phase with respect to the true distribution prevailing in that phase. As the holding times become very large, K tends to zero; but, for phases of finite duration, the (Shannon) information is always less than the prior uncertainty represented by m.e.r.$\{Y\}$. Thus, the data can never suffice to determine the side information in the sense of an error-free demultiplexing operation.

The implication which this bears for the design of adaptive communications receivers is that the wrong subreceiver will be selected at least some of the time in best case. In the "worst case," incorrect decisions by the sdmux will select the wrong subreceiver so often that the performance of the system degrades to below the level of a simpler fixed system. The designer needs some assurance that his adaptive receiver will attain fidelity closer to the best case. Such assurance can be possible only when the sdmux (or a software model of it) has been tested against long records of the same type of composite interference that the actual receiver will be called upon to digest.

Let $\{X\} = \{X(1), \ldots, X(N)\}$ as before, where $X(n)$ is the quantized interference sample at the n-th instant. For N sufficiently large, the number $n(x)$ of observations at quantum level x, after normalization, converges in probability to the true (long-term average) distribution:

$n(x)/N \rightarrow r(x)$ as $N \rightarrow$ infinity.

Moreover, the stated ratio is asymptotically normal to $r(x)$ with variance proportional to $1/N$. But at any given time, the observations are drawn from a subpopulation with distribution $p(x|Y)$ where Y, the state of the process, identifies which subpopulation. As before the states are indexed by a discrete parameter y which is one of M possible numbers. As the Y-process is a Markov chain with stationary distribution $q(y)$, one must have

$$r(x) = \sum_{y=1}^{M} p(x|y)q(y) \ . \tag{10}$$

The sdmux assigns a Y' to every X and, after a long time, the distribution of Y' is seen converging to $q'(y)$. There are, in the limit of continuous y, infinitely many solutions to the simultaneous equations

$$\sum_{y} p(y',y) = q(y') \text{ and } \sum_{y'} p(y',y) = q(y)$$

wherein, if $p(y',y)$ is the joint density of the state Y and the estimated state Y', the marginal densities of these random variables are identical. Unfortunately, none of these fortuitous cases is likely to be obtained in the opertion of the sdmux. If q' and q were the same distribution, every misclassification in which a state of type A is called type B would be "cancelled" by a misclassification of B as A. The sdmux that accomplishes this trick might be called unbiased. The real sdmux will not be unbiased (as might by illustrated by considering, e.g., the case in which Y is the single unknown parameter in an exponential family of conditional densities). This seems the necessary consequence of Equation (9). Thus, the mixture of conditional densities $p(x|y)$, weighted according to the empirical probability distribution $q'(y)$, will not coincide with the long term average $r(x)$ in general.

Thus one has a joint density $p(y,y')$ of the true and the estimated state; and it implies a stationary distribution

$$p(y'|y) = p(y,y')/q(y)$$

of the estimate conditioned on the fact. The estimate $Y'(n)$ is completely determined by the data $X(1), \ldots, X(n-1)$ up to the present time. The average (Shannon) information that Y' conveys about Y is given by

$$I = \sum_{y',y} p(y,y')[\log p(y,y') - \log q(y) - \log q'(y')] \ . \tag{11}$$

Now, if the classification of y as y' is attended by a cost or penalty $C(y',y)$, which is one element of an MxM cost matrix whose diagonal contains all zeros, the average cost or penalty assessed against the sdmux will be given by

$$\sum_{y',y} C(y',y)p(y'|y)q(y) = EC. \tag{12}$$

Equations (11) and (12) are the basis of rate-distortion theory which answers the question, "What is the minimum information (I) required to have the expected cost (EC) less than a criterion level (d)?" This question, with Equations (11) and (12), is answered by the rate-distortion function $R(d)$ for any criterion level d of interest.[6]

The operation of the sdmux on profuse serial data emitted by an MCS (that has the postulated conditional densities) ultimately illuminates the statistics of the Markov chain. Thus q' is known. In addition, since Equation (10) admits a unique solution, q can be computed. Given a plausible cost matrix, $R(d)$ can be calculated. Selecting D as the greatest cost that can be tolerated, $R(D)$ is found to be the minimum information extraction rate that satisfies the requirement. Returning to Equation (9), the mean rate at which the observations convey side information is given by the difference between m.e.r.$\{Y\}$, which is computable from the transition matrix, and the number K, which is likewise computable from the fully partitioned and classified data stream. If

$$R(D) < \text{m.e.r.}\{Y\} - K \ ,$$

then the required information rate is less than the maximum theoretical rate. In this case, the solution generated by the sdmux is verifiable, as one does not need more information than one has (in principle). But if

$$R(D) > \text{m.e.r.}\{Y\} - K$$

an inconsistency appears as the sdmux (or analyst) has posed a decomposition of the data stream which cannot be verified at the criterion level D.

CHAPTER 5

DARNAFALSKI GAMES

In so far as the information theoretic treatment of the adaptive receiver problem up to this point has become increasingly involved in mathematical considerations, and the underlying problem itself is not commonly addressed in either the academic or the applications literature, it may be worth while to illustrate some of the mathematical notions in a more intuitive context. In the study of mathematical statistics, it often happens that an important idea can be illustrated by an imaginary game of chance. Since the student may be more readily able to intuit the implications of the game than the ramifications of the mathematics to some restricted area of professional practice, the use of games to illustrate the theory makes practical sense.

Let the dynamical equation be

$$Z_{n+1} = AZ_n + X_n$$

where A is constant, $X_n = F_n^{-1}(U_n)$, and $U_1, \ldots, U_n, \ldots$ is a sequence of independent Borel trials.* Hold that $F_n(x)$ can be parameterized by Z_n and $Y_n \varepsilon R_y$. Assert that Z_n influences X_n through F_n alone and not through Y_n:

$$\Pr(Y_{n+1} = y | Y_n, \ldots, Y_1, Z_n, \ldots, Z_1) = \Pr(Y_{n+1} = y | Y_n, \ldots, Y_1)$$

---

*Let $B_1, B_2, \ldots$ be an infinite sequence of independent Bernoulli trials and define $U = .B_1B_2B_3 \ldots$ with the RHS being binary representation of a real number on the unit interval.

Then $F_n(x)$ is the distribution function describing X_n, the gain (or loss) accruing, in the n-th round of a Darnafalski game, to a player who started with Z_1 in property.

Now Y_n may be entirely gratuitous, as when $F_n(x) = F(x|Z_n)$ and $\{Z_n, n = 1, 2, \ldots\}$ is a stationary process, the distribution of Z_n, in the limit $n\to\infty$, being related to the initial condition Z_1 through boundary conditions (e.g., absorbing states) or not at all. Games of this first special kind are treated as random walks.

On the other hand, it could be that $F_n(x) = F(x|Y_n)$ regardless of Z_n, so that $\{Z_n, n = 1, 2, \ldots\}$ is a process with conditionally independent increments. When $\{Y_n, n = 1, 2, \ldots\}$ is itself a stationary random process, the prediction problem is posed by the question, "In light of the data $(X_1, \ldots, X_{n-1})$, what will Y_n be?" The answer to this question should influence the decision to play or quit the n-th round of the game. Insofar as $\{X_n, n = 1, 2, \ldots, n-1\}$ is specified by $\{Z_n, n = 1, \ldots, n\}$, Equation (2) makes it clear that all useful information will have been gleaned from the data when they have been used to specify the sequence of parameters.

The nonzero constant A in the dynamical equation makes Z_n a moving average. The process obtained when A = 1 corresponds to the game in which the player's fortune is monetarized in constant value units. When A > 1, the money draws interest. When A < 1, the money in hand is spent or devalued as the game continues. The emphasis, within the present context, is not on games but on nominally <u>zero-mean observation processes</u>. Note that

$$X_n = Z_{n+1} \text{ when } A = 0$$

whereas

$$X_n = Z_{n+1} - Z_n \text{ when } A = 1 .$$

In other words, the sequence of conditionally independent increments in the fair game with constant value units is itself a sequence of conditionally independent zero-mean random variables.

Consider the following game: In the n-th round, the player bets one dollar on the outcome of the roll of a pair of dice which is represented by X(n) which belongs to the set {2, 3, ... ,12}. Actually, there are M pairs of dice from which the house may choose; and the selection of pair identified by the index y engenders the probability distribution

$$P[X(n) = x|Y(n) = y] = p(x|y) .$$

The house follows the procedure of rolling a given pair of dice some finite number of times and then selecting a different pair. The player, who is not informed about the properties of the dice, has only the data X(1), ..., X(N) to use in assessing the odds on having x in the N+1st round. Let the player win B dollars if he guesses correctly. Assume that the rules governing selection of the dice remain invariant.

Based on the last proviso, the best invariant playing strategy must be to bet that X(n) will be g, where

$$\max_x r(x) = r(g) ,$$

for r(x) the normalized frequency of the outcome x up to the present time. The average gain accruing to the player who adopts this strategy is clearly

$$EC(\text{fixed}) = Br(g) - [1 - r(g)] = (B + 1)r(g) - 1 .$$

The player, however, may sumrise that a given pair of dice typically remains in use for many rounds before a change is effected. Therefore, the identity of the dice in use at a given time may be determined with some reasonable level of confidence based on the outcomes of the last several rounds, except when a change has very recently occurred. The adaptive strategy is to place bets on the sequence G(n) where

$$\max_x p[x|Y(n)] = p[G(n)|Y(n)] .$$

Yet the results of the preceding sections apply to the sequence of conditionally independent observations that constitute the game. The most astute player cannot expect to guess the sequence $Y(n)$ correctly. Instead, he poses the sequential hypothesis $Y'(n)$ which leads him to play the sequence $G'(n)$ figured in accordance with the last equation. The expected gain from this strategy is given by an expression analogous to Equation (12) wherein

$$C(y',y) = Bp[G(y')|y] - \{1 - p[G(y')|y]\}$$

for $G'(n) = G[Y'(n)]$. (I.e., $G(y')$ is the largest of the probabilities $p(x|y')$ for given y'.) Substituting this in Equation (12) and rearranging, one has

$$EC(\text{adaptive}) = \sum_{y} q(y)\,\{Bp[G(y)|y]q(y|y) -$$

$$\sum_{y' \neq y} p[G(y')|y]q(y'|y)\}$$

when $q(y'|y)$ is understood as the distribution of y' conditioned on y. Now the convention attached to the cost matrix in the communications context was that its diagonal vanishes and its other elements are non-negative; so the cost is always positive. In calculating the expected gains from the fixed and adaptive playing strategies, we have positive numbers on the main diagonal and negative numbers elsewhere; and there is no assurance that the expected gain is non-negative. Indeed, if B is too small, the optimum fixed and adaptive strategies may both be losing strategies. Despite these technical difficulties, it would seem plausible to suppose that the rate-distortion function exists and can be computed by methods not much different than the standards. More precisely, if the information is defined as

$$I = \sum_{y',y} q(y',y)[\log q(y'|y) - \log q'(y')],$$

there would seem to be no serious impediment to the computation of

$$R(D) = \min(I)$$

such that EC(adaptive) is not less than D. The minimization extends over all conditional distributions $q(y'|y)$. If

$$R[EC(\text{fixed})] = i$$

defines the minimum information necessary for the adaptive strategy to yield the same expected gain as the optimal strategy to yield the same expected gain as the optimal fixed strategy, then the player needs to know whether the information contained in the observations is greater than i. For if the sequence of rolls does not convey information about the nature of the dice at a mean rate more than i, then the adaptive strategy must be inferior. On the other hand, if the theoretical information rate computed in accordance with Equation (9) is more than i, the player has the assurance that the adaptive strategy can be made to work better. Bear in mind that the question of how to implement the adaptive strategy (in the "real time" situation of the game player) remains thus far unresolved. The next chapter presents some mathematical background material which is relevant to this question.

CHAPTER 6

MEASURES OF STOCHASTIC DIVERGENCE

Consider two sequences of iid random variables, X(1), ..., X(N) and X'(1), ..., X'(N'). Let X have the distribution function F(x) and let X' have the distribution function F'(x). If F = F' at every x, then the two sequences are said to be stochastically equal. If F(x) is greater than or equal to F'(x), then X is stochastically less than X'. Similarly, if F(x) < F'(x), then X is stochastically greater than X'.[7]

Equation (4), which defined the joint density of the conditionally independent observations X(n) given the parallel sequence {Y(n)} of unobservable states, suggests that the observations can be sorted into subpopulations using some criterion for measuring stochastic equality *to* any of a number of conditional densities $p(x|y)$. Obviously, the block of data X(1), ..., X(n), if all the individuals are drawn from the same subpopulation identified by a particular y, defines a sample distribution (or empirical distribution) function, denoted F(x;n), to which there corresponds a sample density, denoted f(x;n), which should look very much like $p(x|y)$ for sufficiently large n. There are a number of well known techniques for comparing f(x;n) to $p(x|y)$ and passing judgment on the so-called goodness-of-fit. Each technique begins with a formula for measuring the divergence of the sample distribution (or density) from the theoretical one.

<u>The integrated squared difference</u> is[8]

$$J(n,y) = \sum_{x} [f(x;n) - p(x|y)]^2 .$$

The variation is[9]

$$V(n,y) = \sum_{x} \left| f(x;n) - p(x|y) \right| .$$

The Kolmogorov distance between the two distributions is defined by

$$D(n,y) = \sup \left| \int_{-\infty}^{x'} [f(x;n) - p(x|y)]dx \right|$$

when x is absolutely or piecewise continuous.

The Kullback directed divergence is

$$K(n,y) = I(f;p) = \int_{-\infty}^{\infty} f(x;n)\log[f(x;n)/p(x|y)]dx$$

and it is defined when the argument of logarithm has no zeros in the denominator. The other directed divergence implied by this definition is $I(p;f)$; and the sum $I(p;f) + I(f;p)$ is the undirected divergence of p and f.[10]

Each of these measures of stochastic divergence tends to zero as the size n of the sample goes to infinity under the assumption that the samples are indeed drawn from the subpopulation indicated by the parameter y. Otherwise, they converge to positive values which represent the distance, divergence, or disparity between the postulated distribution and the true distribution generating the observations in the sample. The statistician will want to know the distribution of the divergence statistic under the null hypothesis (of stochastic equality), preferably for any n, but at least for large n. In order for there to be a single distribution that answers the question, the test of stochastic equality implied by the divergence statistic must be a distribution-free test. Kolmogorov's statistic yields a distribution-free test and it has been tabulated for small n.[11] The Kullback undirected divergence is asymptotically distribution-free.

The Kullback directed divergence is of particular interest in the present context owing to Equations (8) and (9) through which it figures in the computation of the rate at which the observations convey side information. Equation (8) contains Q(k), the fraction of the total time which passes in the k-th phase. As the data stream continues ad infinitum, every Q(k) goes to zero; but the sum over all phases of those Q(k) for which the process is in state y approaches q(y), the stationary probability of finding the process in state y. Therefore, Equation (8) is the same as

$$K = \sum_{y} \sum_{n=1}^{\infty} q(y)u(n|y)K(n,y) \tag{13}$$

where u(n|y) is the distribution of the holding time of the Y-process in state y and K(n,y) is the Kullback directed divergence defined above. This makes explicit the earlier assertion that K is the average Kullback divergence of the homogeneous subsequence from the parent distribution.

This result suggests a modus operandi for the sdmux. As before, the practical situations of interest involve holding times which are typically large. The mean holding time of the process in the y-state is

$$\bar{n}(y) = \sum_{n} nu(n|y) \; ;$$

and the mean holding time of the process is

$$\bar{n} = \sum_{y} \bar{n}(y)q(y).$$

If $\bar{n}(y)$ is at least several times larger than some integer B for every y, then partitioning the string of N observations into consecutive blocks of B observations will produce blocks which are often homogeneous in stochastic type. Some of the blocks will contain observations of two or more types; but many of these will be dominated by a single type with only a small number of contaminants. Of course, it is presumed in making these statements that B itself is a "large" number.

Let the sdmux operate on the observations by taking them in blocks of length B and computing K(B,y) for each possible y. One particular state, call it Y', has the property

$$K(B,Y') = \min K(B,y) \tag{14}$$

when the minimum over all y is selected. Then, Equation (14) defines the state of the process in this block of observations according to the sdmux. Suppose that the sdmux classifies all the consecutive blocks of data on this basis and tacitly admit that its classifications all are correct. Then the average prescribed by Equation (13) will be the same as the average value of K(B,y'). By operating in the suggested manner and averaging its selector statistic K(B,Y') as it goes along, the sdmux computes the difference between the mean (Shannon) information rate and the mean entropy rate of the side process as shown in Equation (9).

It can be shown that the distribution of BK(B,Y') is asymptotically chi-square with J/2+1 degrees of freedom where J is the number of elements in the sample space of the observation. The proof is somewhat unconvincing; but Monte Carlo calculations performed for the author show that the theorem is plausible for a variety of underlying distributions. Therefore, the average value of the selector statistic, when the underlying model is perfectly valid, is on the order of (J/2+1)/B when the block in question is truly homogeneous. When a block spans the boundary between two phases, larger values result. Thus

$$K > (J/2+1)/B$$

gives the lower bound. Since block length is here synonymous with sample size, it is not surprising that B must far exceed the cardinality of the range of X for the method to succeed. Hence, the proviso that B is "large" translates more specifically to the requirement that it be large compared to J.

Selection of the proper block length (B) would now appear to be a critical aspect of sdmux design. As B grows longer, the theoretical information extraction rate increases; but at some point B becomes so large that it typically spans more than one homogeneous segment of the data stream. Thereupon K begins to diverge and further increase in the block length will degrade the fidelity of the demultiplexing operation.

CHAPTER 7

DATA SCREENS AND THE INFORMATION INDICATOR

In a great variety of R&D tasks, the need arises for methods to identify the more significant portions of a large body of raw data in order to subject them to detailed analysis. Frequently, the analytic procedure will have been established beforehand; but because it is cumbersome, tedious, and requires special talents or instruments, to process all of the available data is not practical. Yet although the procedure for extracting the information from the data is complex, the identification of those parts of the whole record which contain the information may be simpler. Accordingly, a data screening procedure is defined by a set of rules for the purpose of discarding the insignificant data in order to expend a larger share of the total time (or energy) on the anaysis of the remainder. The rules defining the screen may be rather loose, as when the chief scientist relies on trained assistants to take a quick look at the whole data record and then submit the more interesting portions to him for indepth analysis. On the other hand, if the data can be conveniently read into a computer which can extract all the relevant information through the application of well-defined mathematical methods, the screening procedure may be applied to the reduced data, or the need for it obviated altogether. This could be the general case when laboratory analysis of data is considered, as laboratory analysis is constrained to draw conclusions by a date far subsequent to the acquisition of the data.

The situation is different with regard to real time data processing applications. When the data are analyzed in real time, the information received at a certain instant must be extracted before a fixed length of time has elapsed. This lag is typically on the order of seconds. Updating applications will here refer to cases in which the data processor falls behind intermittently but always catches up, as the average length of the information backlog does not

increase with time. In these situations, speed is the essential problem. Although the trend towards increasing computing power is ever smaller and more efficient packages may continue at the astounding pace demonstrated in recent years, designs for automatic control and artificial intelligence will inevitably desire more than can be obtained off the shelf. Consequently, there will be a continuing need for economy in programming as proven algorithms for laboratory data analysis fail to fit the requirements of real time and updating applications.

A distinct class of problems in the realm of updating applications involves data which are received continuously but in which the information content is sporadic. The receiver system relies on a computer to extract the information from the data; and if the computer is dedicated to this task, it will output nothing significant for much of the time. Imagine a multichannel receiver system with the updating requirement that uses a computer too slow to analyze all of the incoming data, but sufficiently fast to extract all relevant information if fed only the information-bearing segments. The receiver is designed so that each channel is screened and a buffer holds the backlog of data awaiting processing by the computer. The data screens are applied to all the channels continuously and the unscreened (remainder) portions are placed in a single queue to be considered sequentially. Let the channels be labelled with the index n and define I_n the average information rate on the n-th channel. The combined average information rate at the input of the processor is

$$I = \sum_n I_n \tag{14}$$

which must be exceeded by the rate R of information extraction. Now the channels may be qualitatively different and the number of computations required to extract J bits from the n-th channel may be different from the number of computations to extract J bits from the m-th channel. Hence the information extraction rate R_n is channel specific. Then

$$R = \sum_n A_n R_n \tag{15}$$

is the average information extraction rate, weighted according to the channel information rates, where $A_n = I_n/I$. The backlog will vanish intermittently if R is greater than I. If I exceeds R, the backlog will grow with the passage of time until the buffer overflows (no matter what its capacity). Using equations (14) and (15), it can be shown that $R > I$ implies

$$\sum_n I_n R_n > (\sum_n I_n)^2$$

which is trivially solved when every $R_n > I$. (For the single channel case this is the tautological solution.) A more thorough examination using some kind of stochastic queueing model would show how, for example, the mean length of the backlog declines as the data processing rate increases.

A second distinct class of updating applications involves isolated receiver systems, i.e., those which draw power from a limited reservoir of energy. Suppose the information rate on a single channel is I and the processing rate is $R > I$. Then the data will be placed in a buffer or storage register until J bytes have been collected, at which time the stored bytes will be read into the processor. When data are being collected in the buffer, the processor has nothing to do. During a long time interval (0,t), the processor will be working for t_1 seconds and idle for t_o seconds, with $t_o + t_1 = t$. The equality of information bytes into and out of the processor implies

$$Rt_1 = It \ . \tag{16}$$

If the processor is free running, it draws power at a rate P_1 (watts) whether it is handling information or not. (This is true of NMOS and bipolar devices.) Its total power requirement then has the steady value

$$P = P_1 + P_o$$

where P_o is the power drawn by all the other system elements together. But the processor could be gated ON when the buffer is full and OFF when the J bytes have been processed. In this case the required energy up to time t is

$$W(t) = P_1 t_1 + P_o t \ .$$

The average power required by this gated, buffered system is

$$\bar{P} = W(t)/t = P_1 I/R + P_o \qquad (17)$$

with reference to Equation (16).

Figure 1 shows $\bar{P}(R)$ in comparison with the power P of the free running system. Obviously, service life can be extended by gating a fast but power-hungry processor in an isolated receiver system. Note that if an isolated multichannel system were considered, the result would be the same, with R given by Equation (15), whether the information rates are steady or sporadic. In fact, the curve in Figure 1 also represents the average power drawn by a system which receives information sporadically and uses a screen to reject the superfluous data, turning on the information processor only when information is detected. When interpreted in this context, Equation (17) carries the proviso that the screening procedure is infallible, since any unnecessary awakening of the processor carries an energy penalty.

False activations of the processor may be rare if one has, e.g., a noise-free digital channel. Then the receipt of a few start bits is easily detected by a single chip device which enables loading of the subsequent data burst into a processor or recorder. Noise will generate spurious bits, triggering false activations of the processor or loading unintelligibles into the recorder. This could be a serious problem if the processor is programmed to fill its memory with a certain number of bits at each activation. The obvious solution is to screen the data by rejecting bursts of less than the right number of bits. The screening device would have two parts. First is a buffer which holds a number of bits while their information content is evaluated. Second is a logic array to perform the evaluation. In this specific case, the logic array might just be an adder with a threshold test. More generally, the logic array could be called an Information Indicator (II), since its function is to output logic one when it detects information, and logic zero otherwise. The data

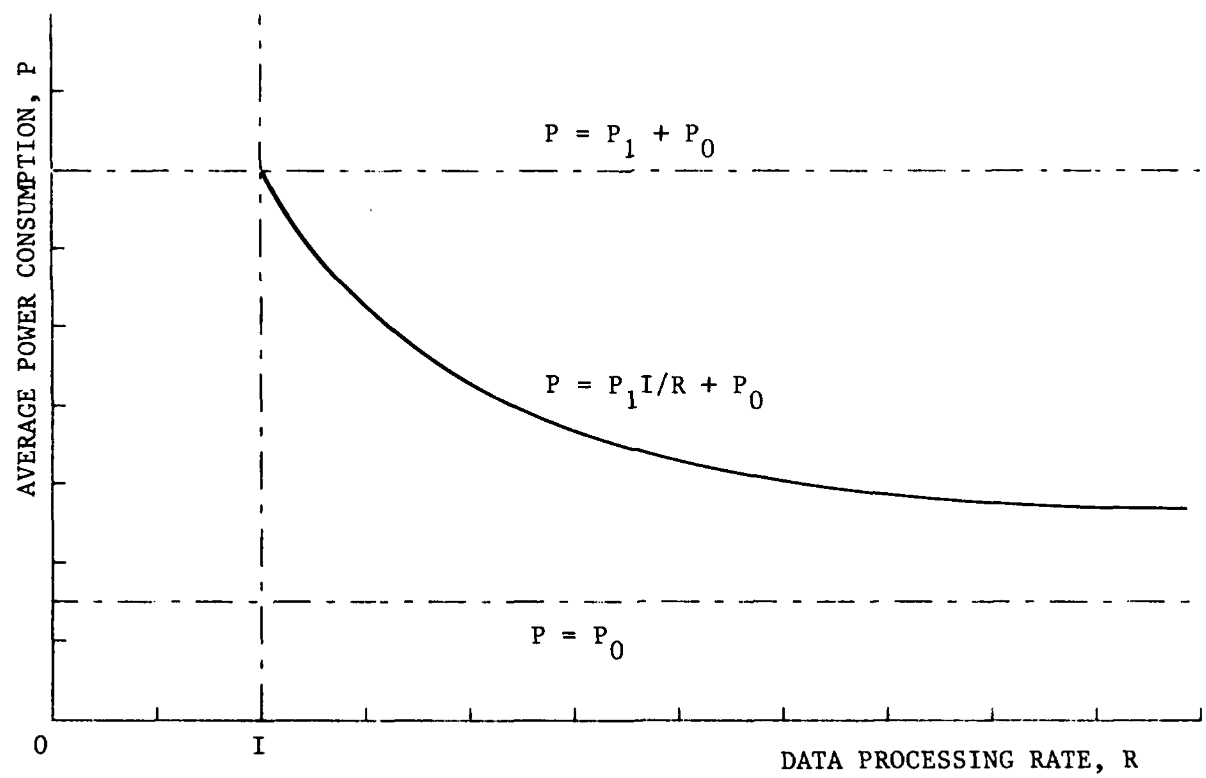

FIGURE 1. AVERAGE POWER CONSUMPTION VERSUS DATA PROCESSING RATE FOR THE GATED AND BUFFERED PROCESSOR

screen, realized in terms of digital (or analog) hardware, is then an Information Indicator/Buffer (II/B). This study will propose some general mathematical ideas for the design of II/B's for isolated receiver systems.

The advantage of using an II/B in a particular application will depend on information rates, data processing rates, and device power requirements, as illustrated in Figure 1; and also on the accuracy of the II/B. Before considering the question of accuracy (in terms of false activation and false rest probabilities), it is worthwhile to evaluate some hard data on device power requirements. The utility of the II/B is predicated on the assumptions of a power-hungry processor and a screening algorithm that is much more efficiently realized in terms of hardware. The first assumption is readily tested with reference to some manufacturers' data on the present generation of microprocessors. Table 1 lists the clock rates and power requirements of eight devices, numbered on the left beginning with the most efficient. Naturally, the CMOS take less power than the NMOS devices; and the one bipolar entry requires the most power in addition to using the fastest clock. But speed is determined by microprocessor architecture and does not necessarily correlate well with clock rate. Suppose the data processing algorithm to be performed by the microprocessor in the isolated receiver system involves a significant number of 16-bit multiplications. Using the same algorithm for doing multiplications and implementing it with the instructions used by the particular device, an unbiased consultant has arrived at the numbers of clock cycles (per multiplication) listed in column #6 of Table 1. Here a tremendous variance exists. The data in column #6 together with columns #4 and #5, imply columns #7 and #8 giving respectively the millijoules and microseconds per 16-bit multiplication for each of the eight devices. In addition, devices 9 and 10 are single-purpose multiplier chips, listed for comparison. The speed and energy data are displayed graphically in Figure 2. Remarkably, the power-hungry bipolar microprocessor (row #8) performs 16-bit multiplications using about the same amount of energy as the CMOS devices (rows #1 and #2), while working at a much faster rate. Incidentally, the device in question (AMD 2903) is built on four separate chips in order to dissipate the heat (7.0 watts) produced in the free running mode.

TABLE 1. MICROPROCESSOR DATA

1	2	3	4	5	6	7	8	9	10
RANK	DEVICE	CLOCK RATE f_C (MHz)	POWER P_D (mW)	P_D/f_C (nJ)	CYCLES PER MULTIPLICATION (16-BIT NOS.)	EXECUTION TIME t_m (μsec)	EXECUTION ENERGY (mJ)	DEVICE TYPE	MANUFACTURER(S)
1	1802	4 (max.)	$4f_C$	4	2560	663	10	CMOS	RCA
2	Z80	2.5	50	20	2272	909	45	p^2CMOS	NSC
3	MC14-6805	1	20	20	504	504	10	CMOS	Motorola
4	6502	2	200	100	464	232	50	NMOS	Commodore, Rockwell
5	9511	4	850	210	90	23	20	NMOS	AMD
6	6800	2	600	140	464	232	140	NMOS	Motorola
7	8080A	2	780	590	1504	750	590	NMOS	INTEL
8	2903	10	7000	700	17	1.7	12	bipolar	AMD
9*	25LS14	--	--	--	--	1.3	1.4		AMD
10*	CDP1855	--	--	--	--	0.7	0.14	CMOS	RCA

*indicates device is a single-purpose multiplier chip

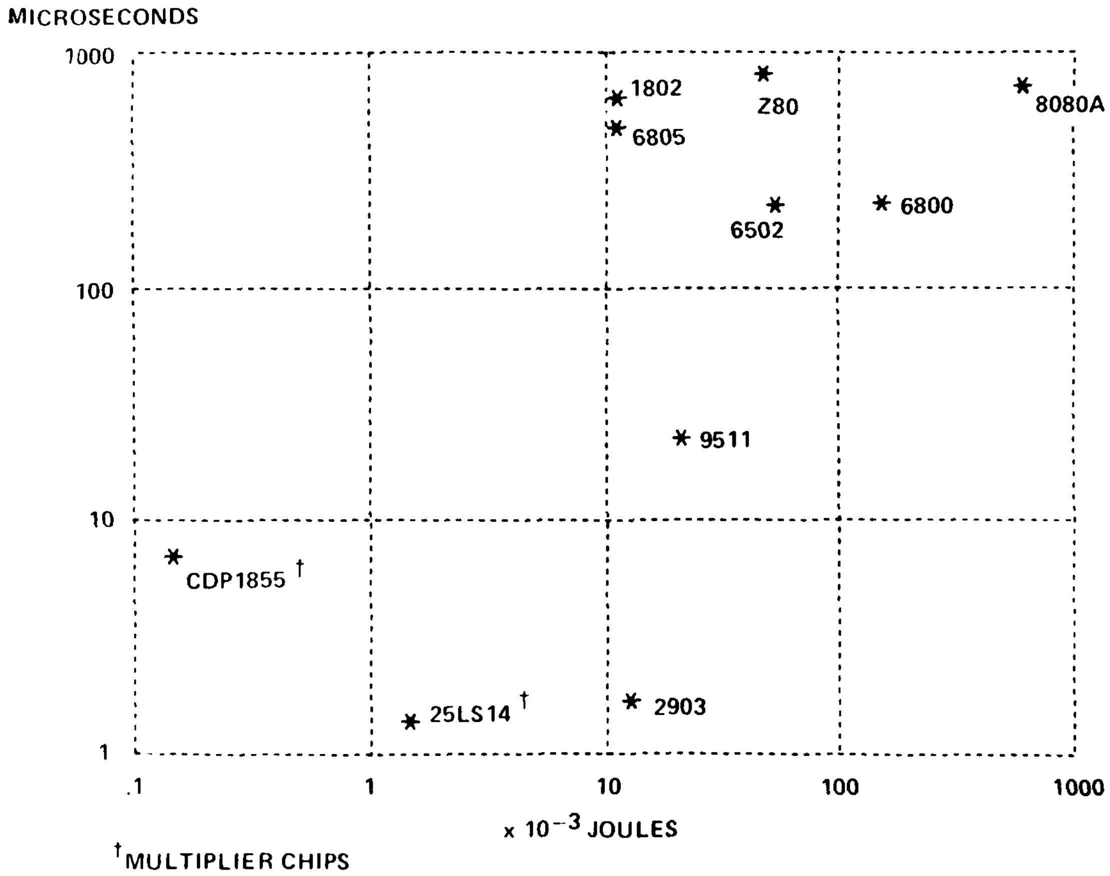

FIGURE 2. EXECUTION TIME VERSUS ENERGY

The rationale for using the time and energy per multiplication as yardsticks to measure processor performance lies in the fact that matrix inversion requires a large number of multiplication operations; and the theory of estimation relies heavily on matrix inversion techniques. If the task of the processor is to estimate a scalar X based on noisy observations Y_{-1}, Y_{-2}, ..., Y_{-n}, then it finds the best linear estimate

$$X' = A_{XY} Y A_{YY}^{-1} , \tag{18}$$

where Y is the observation vector and A_{XY} and A_{YY} are nxn covariance matrices. Just the number of multiplications required to invert A_{YY} is

$$M(n) = n^3 + n(n - 1)^3 , \tag{19}$$

although it may be reduced when some elements are identically zero. Even the Kalman filter, praised in some texts for its ease of implementation on computers because it computes the error covariance matrix recursively from the preceding value, inverts the matrix every time it computes an estimate. Now if the processor must cycle through the estimator/filter routine each time it receives a digitized sample of a noisy signal, and the sampling frequency is f_o, then the inequality

$$f_o < 1/M(n)t_m \tag{20}$$

must be satisfied to permit real time operation, where t_m is the time required to do a multiplication. Thus when t_m is fixed by selection of a particular device, the number of elements in the observation vector must be fewer than

$$\max(n): \; M(n) < 1/f_o t_m , \tag{21}$$

Equations (19) and (21) were used to prepare the chart in Figure 3 for the values of t_m corresponding to the devices with least quiescent consumption (row #1) and greatest speed (row #8).

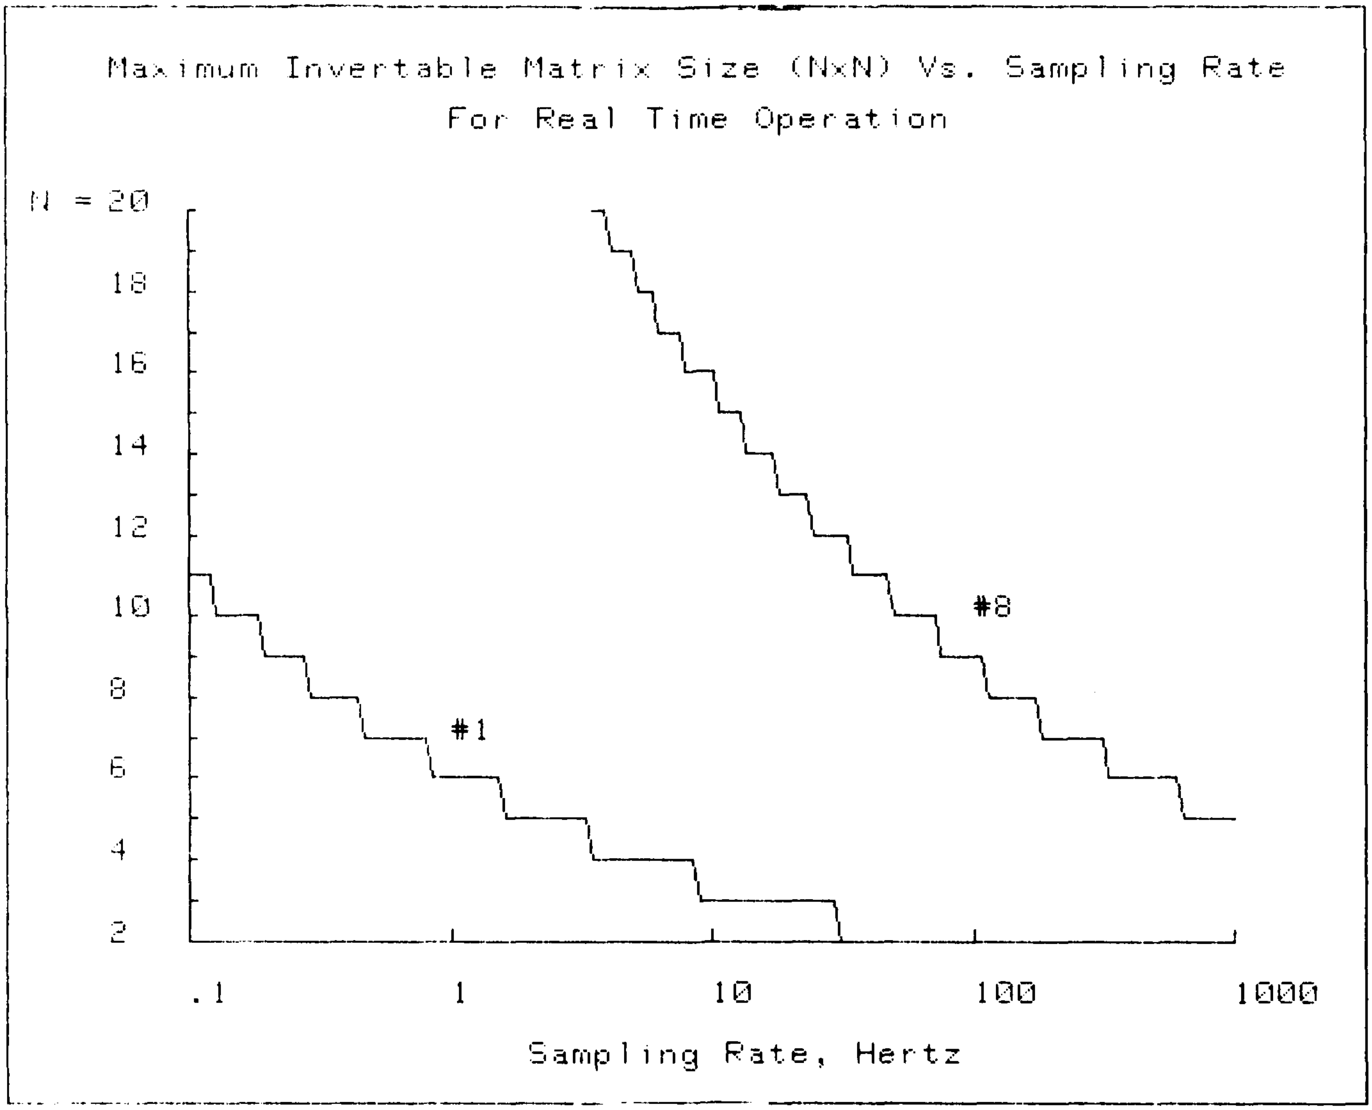

FIGURE 3. MAXIMUM INVERTABLE MATRIX SIZE (NxN) VERSUS SAMPLING RATE FOR REAL TIME OPERATION

In summary, the manufacturer's data on the present generation of microprocessor devices show that the two goals of high data processing rate and low quiescent power consumption have been achieved by separate groups of designers. The comparison of data processing devices in the same class and generation might well show that the most efficient devices are not the fastest, and vice versa, when efficiency is judged on the basis of free running power consumption. A better measure of efficiency is the reciprocal energy required in a standard complex operation. Taken in this latter sense, maximum efficiency and speed might be achieved by using an ostensibly power-hungry device which is gated on for brief intervals according to the information content in a buffer. A drawback to this approach is that the processor and II/B must be entirely divorced. With a less voracious processor, the II/B algorithm and the actual data processing could be done by the same device. When the output of the II routine changes to logic one, the program branches into the data processing mode, perhaps shifting the clock to a faster rate at the same time. But the use of a power-hungry processor to extract information from the data dictates use of separate package of II/B hardware. Depending on the complexity of the II algorithm, this separate package may be a second microprocessor that features very low quiescent power consumption.

CHAPTER 8

ALERTNESS STRATEGIES FOR ISOLATED RECEIVERS

When information comes to the receiver intermittantly via the additive noise channel, the device which performs the information indicator function might be called a signal detector. A Simplified Signal Detector (SSD) is proposed for use in conjunction with computationally advanced but power-hungry signal processors in isolated receiver systems where energy conservation will extend service life. Figure 4 is a diagram of the receiver system. A sensor produces a voltage x(t) which is passed through a filter. The filter may be designed for anti-aliasing, whitening, or any combination of purposes; but it is assumed to be a time-invariant network. The filter output is y(t). The sensor and filter together draw I_A amperes from the supply voltage V_S. The SSD examines y(t) and makes a decision every T seconds as to whether the sensor is reporting noise only (hypothesis H0) or noise combined with some type of signal (H1). If H0 is selected, the output z of the SSD remains low (logic zero). If H1 is selected, z goes high (logic one) and closes the switch SQ, activating the advanced processor. The currents drawn by the SSD and the processor are I_B and I_C, respectively. Once activated, the processor will run for t_o seconds, decoding the signal if indeed a signal is present. The processor can override the SSD and sustain its own connection to the power source, after it has locked onto a message. It is assumed that H1 is a rare event; i.e., the system waits long times between signals. Suppose that the power source is comprised of batteries so that service life is determined by current drain.

The SSD can commit errors of two kinds--false activation ("H1"|H0) and false rest ("H0"|H1). Let

$$Q_o = \text{Prob}(\text{"H1"}|\text{H0})$$

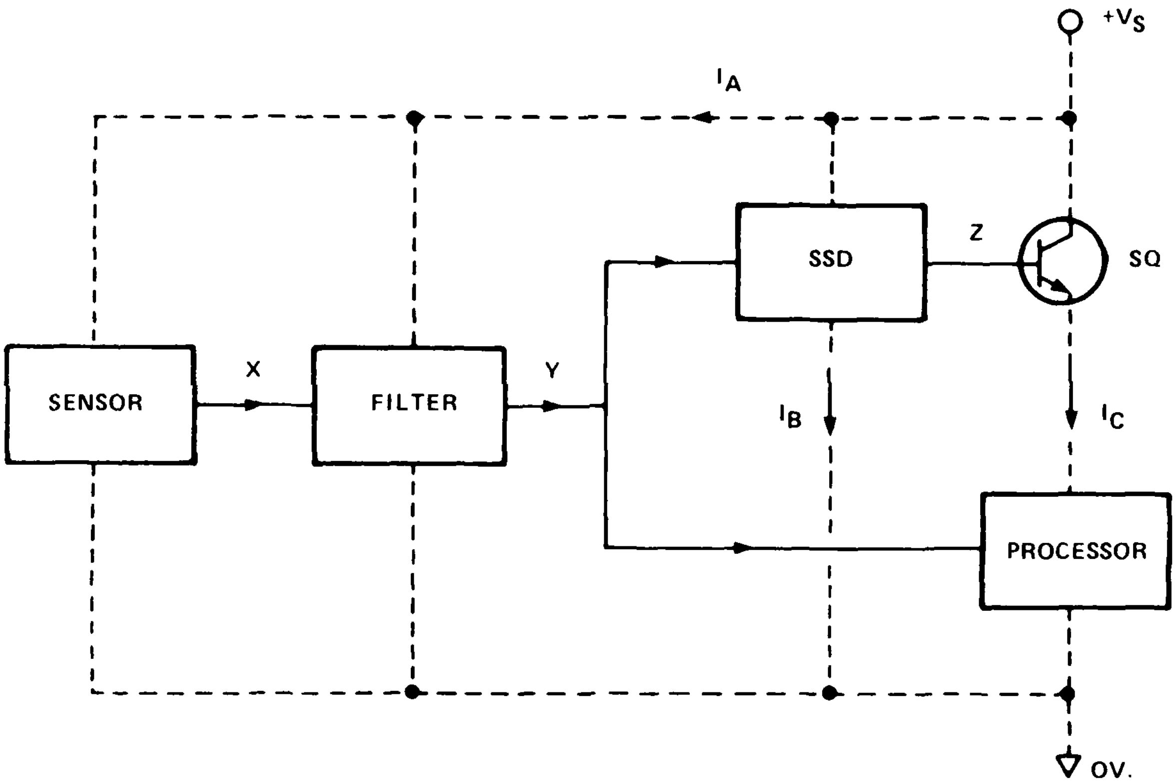

FIGURE 4. ISOLATED RECEIVER SYSTEM USING SIMPLIFIED SIGNAL DETECTOR

be the probability of an error of the first kind. The SSD samples y(t) every T seconds for a sampling rate of $f_o = 1/T$. Then the rate of occurrence of false activations is $Q_o f_o$ on the average. If the system is started up at time zero, then at a later time t the expected number of false activations is

$$N = Q_o f_o (t - N t_o) \tag{22}$$

since the processor was already activated during $N t_o < t$ seconds. Equation (22) is the same as

$$N = Q_o f_o t/(1 + Q_o f_o t_o) \tag{23}$$

Note that N approaches $Q_o f_o t$ as the mean number $Q_o f_o t_o$ of errors per activation interval (t_o) approaches zero.

The power supply capacity (in ampere-seconds) consumed by the processor up to time t is $N t_o I_C$ under the assumption of no overrides. I.e., it is assumed that every activation of the processor up to time t has been triggered by an error of the SSD. Then the current drawn by the system has the time-average value

$$Av(I_1) = I_A + I_B + N I_C t_o = I_A + I_B + I_C Q_o f_o t_o/(1 + Q_o f_o t_o) \tag{24}$$

with reference to Equation (23). Elimination of the SSD from the design puts I_B=0 while holding SQ in the conducting state. The average current drain in this baseline (free running) system is

$$Av(I_o) = I_A + I_C \ . \tag{25}$$

Since the service life is inversely proportional to current drain, the SSD extends service life by a factor

$$L = Av(I_o)/Av(I_1) \ . \tag{26}$$

Substituting Equations (24) and (25) in Equation (26) with the additional definitions $A = I_A/I_C$ and $B = I_B/I_C$ gives

$$L = \frac{A + 1}{A + B + \dfrac{Q_o f_o t_o}{1 + Q_o f_o t_o}} . \qquad (27)$$

Now define

$$G = Q_o f_o t_o = \text{mean errors per activation interval.}$$

Rearrangement of Equation (27) immediately shows

$$G = \frac{1 + (1 - L)A - LB}{L - 1 -(1 - L)A + LB} , \qquad (28)$$

the error rate required to achieve a given life extension factor for circuit parameters A and B. Figure 5 shows L(G), the life extension factor versus error rate, for several values of A = B. Figures 6 and 7 plot the same function for A = B/3 and A = 3B, respectively.

When the task of the SSD is to discern a weak signal, ie., to operate in a low signal-to-noise ratio (SNR) environment, the formulation of the II function may be challenging. Such formulation requires a theoretical solution of the binary decision problem and the means to realize the solution with hardware. The designer faces a tradeoff, having to compromise the objectives of long service life and high detection probability. The SSD will reduce the overall performance of the system by failing to detect the signal with probability $Q_1 = 1 - Q_d$, where Q_d is the probability of selecting H1 when H1 is true.

Determination of the optimum decision threshold is strongly influenced by the costs or benefits assigned to the four contingencies ("Hi"|Hj, i,j = 0,1). A special case is considered here. The advent of the signal is a rare but recurrent event; and over a long period of time the number of messages reflects an average rate of occurrence. In this case, the receiver system will interpret a number of messages proportional to the product of its lifetime and the

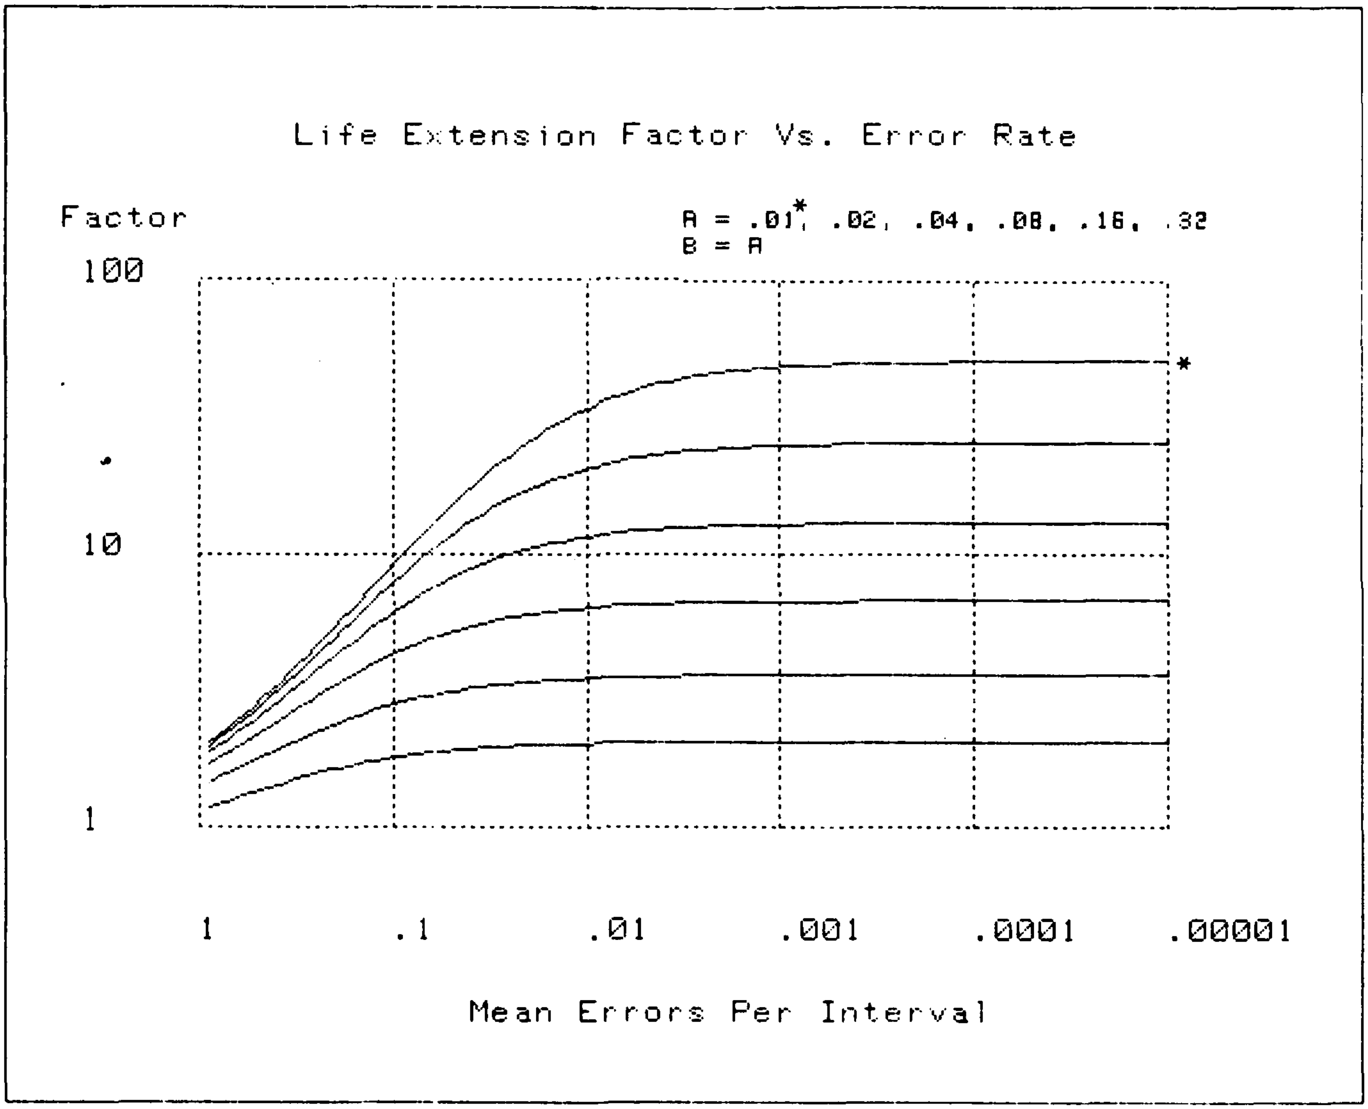

FIGURE 5. LIFE EXTENSION FACTOR VERSUS ERROR RATE

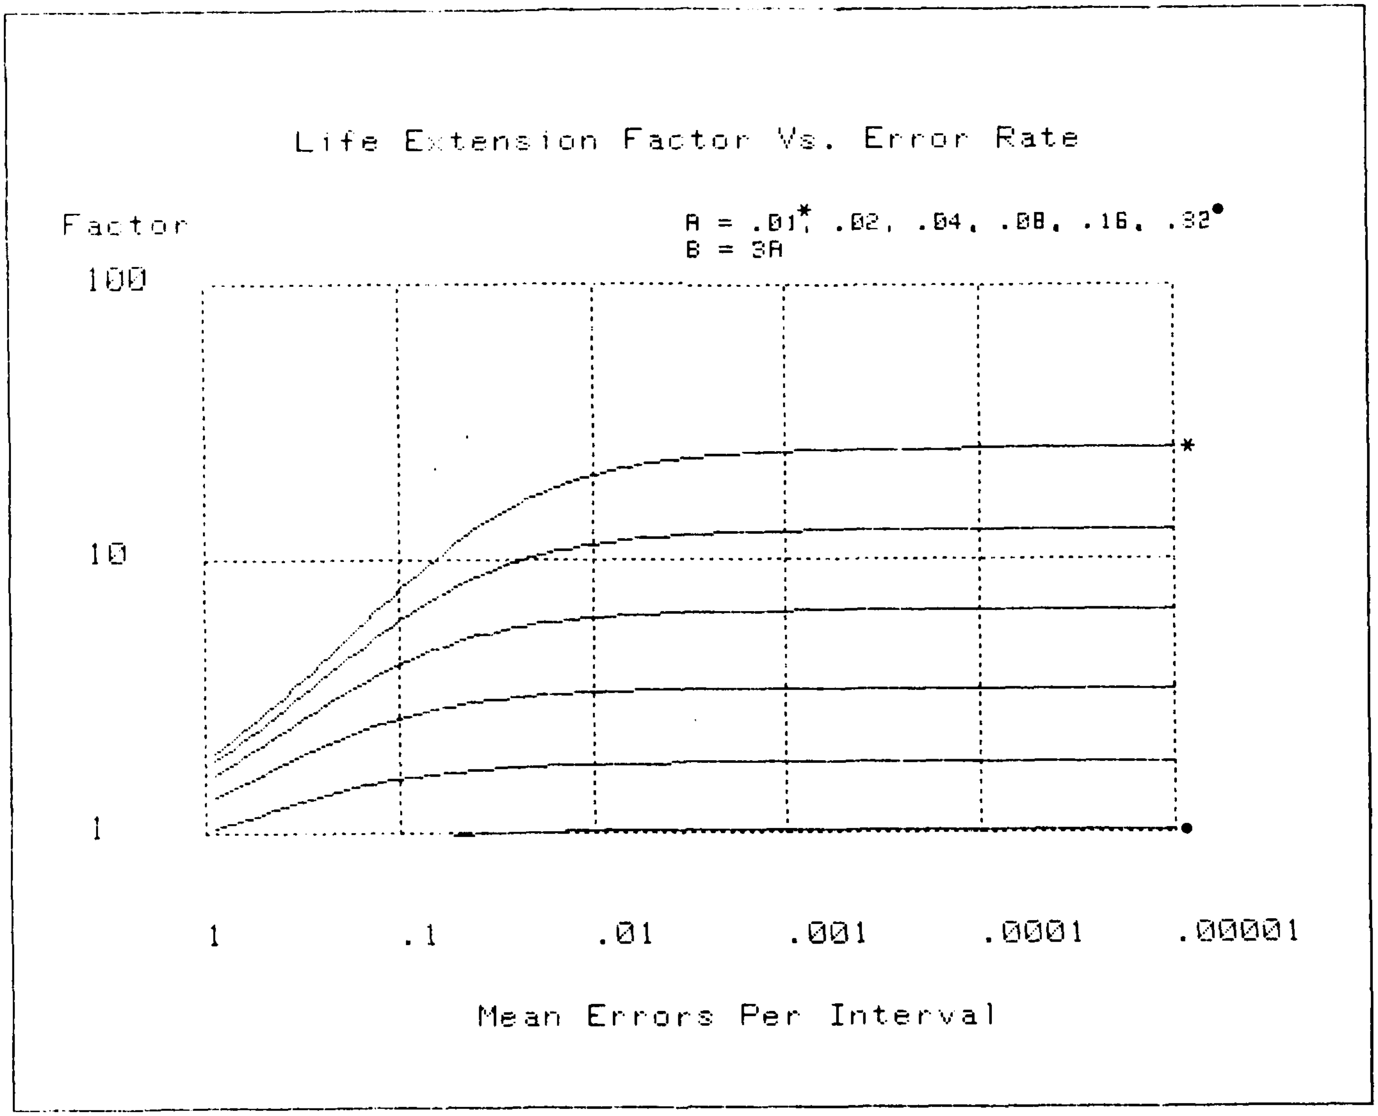

FIGURE 6. LIFE EXTENSION FACTOR VERSUS ERROR RATE

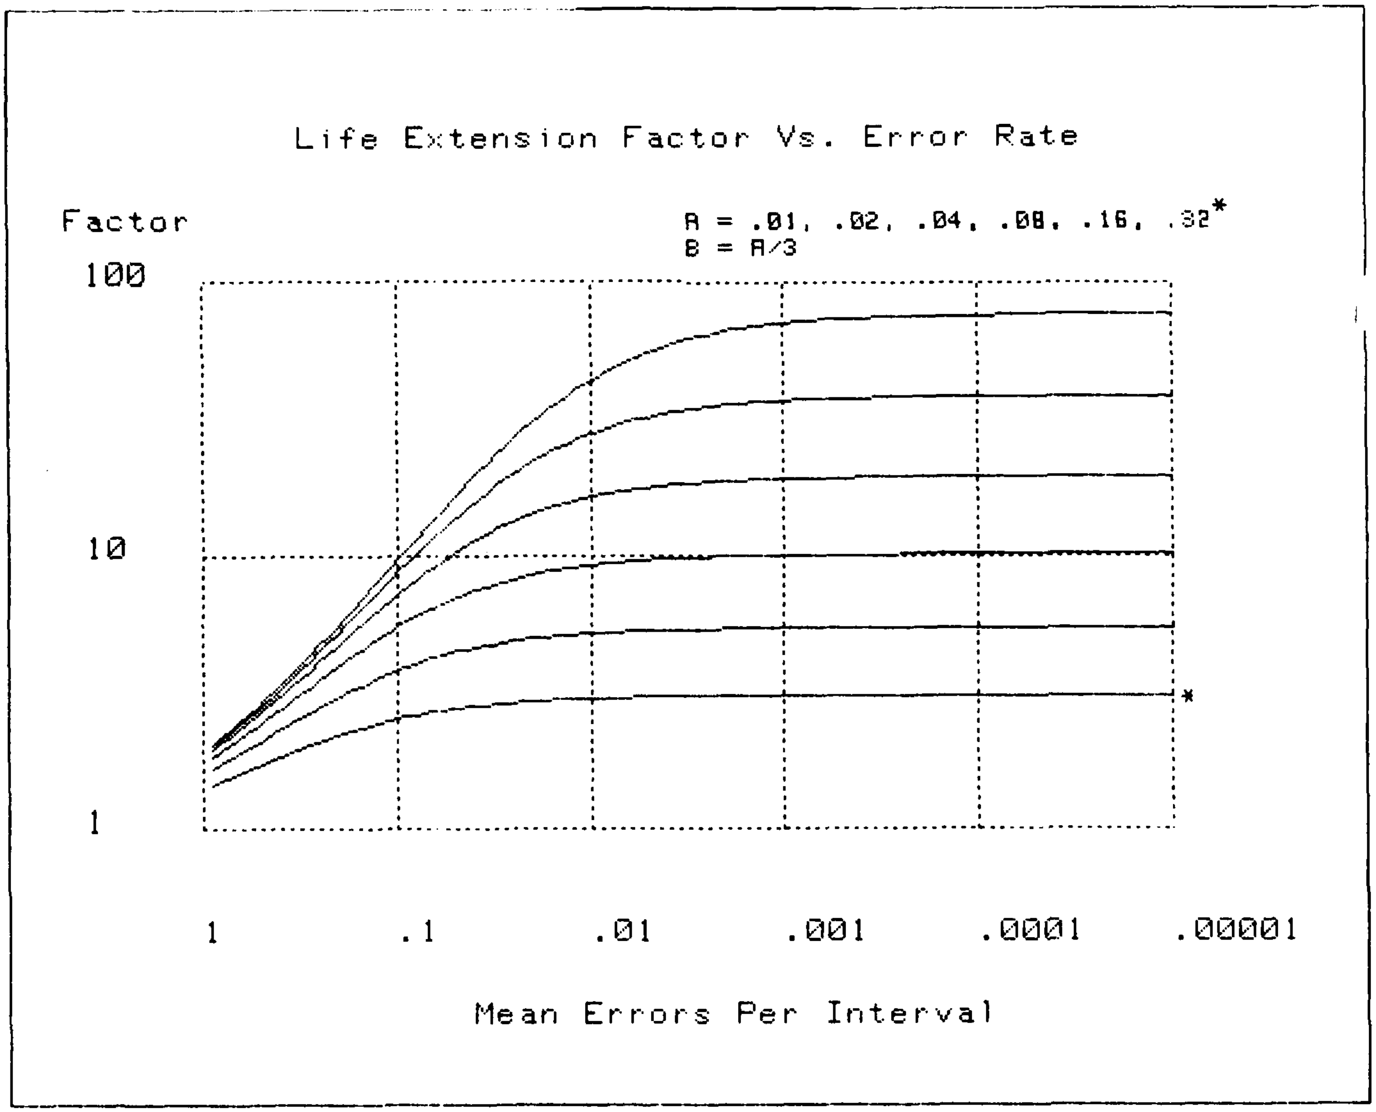

FIGURE 7. LIFE EXTENSION FACTOR VERSUS ERROR RATE

detection probability of the SSD. Thus if the baseline receiver system, which lacks an SSD, has unit life expectency and 100 percent probability of correctly interpreting every message, its value is one. The extended-life system with the SSD has value

$$V = L(Q_o)Q_d \tag{29}$$

in comparison, where (Q_o,Q_d) is the operating point of the SSD and $L(Q_o)$ is given by Equation (27) or some other expression. Since Q_o and Q_d are functions of the decision threshold z', the value of the system is a function V(z') of the threshold setting.

Since $L(Q_o)$ is a decreasing function, and $Q_d(Q_o)$ is increasing, it may be expected that V(z') has a unique maximum for some optimum threshold z" for which the false activation probability is $Q_o(z")$. Chain rule differentiation of Equation (29) yields

$$dV/dz' = (dL/dQ_o)(dQ_o/dz')Q_d + LdQ_d/dz' , \tag{30}$$

Before setting this to zero and rearranging, note that

$$(dQ_d/dz')/(dQ_o/dz') = dQ_d/dQ_o = p(z'|H1)/p(z'|H0) \tag{31}$$

which is usually called the likelihood ratio and given the symbol Λ. Then the solution, if it exists, is that threshold which satisfies

$$-\Lambda = (Q_d/L)dL/dQ_o = Q_d d(\ln L)/dQ_o. \tag{32}$$

Specifically substituting the R.H.S. of Equation (27) in Equation (32) gives

$$\Lambda = \frac{Q_d}{A + 1} \frac{f_o t_o L(Q_o)}{(1 + f_o t_o Q_o)^2}, \tag{33}$$

the solution of which requires some computation in any nontrivial case. When the optimum threshold is found to correspond to a detection probability which is nearly one, and the ratio of currents A is much less than one, Equation (33) says

$$A \doteq M_p L/(1 + G)^2 \text{ for } z' = z'' \tag{34}$$

approximately, where $M_p = f_o t_o$ is the number of data samples taken up by the processor after the SSD gates it ON. It has been shown that the life extension factor is large only when the number G of errors per activation interval is also small. Hence

$$A \approx M_p L \tag{35}$$

is the implication.

To illustrate Equation (35), let the test statistic be the normalized sum of squares of the past N inputs:

$$Z = [X^2(1) + \ldots + X^2(N)]/N \,. \tag{36}$$

Then the conditional means of Z are

$$E(Z|H0) = m(0)$$

and

$$E(Z|H1) = m(0) + m(1)$$

where m(1) corresponds to signal power and m(0) to noise power. Since Z/m(0) is chi-square with N degrees of freedom under the null hypothesis,

$$Var(Z|H0) = 2m^2(0)/N \,. \tag{37}$$

When N is large, the chi-square distribution is close to normal and it follows that

$$\log\Lambda = Nm(1)[Z - m(1)/2 - m(0)]/m^2(0) \; ; \tag{38}$$

so the threshold

$$z'' = m(0) + m(1)/2 + \frac{m^2(0)M_p L}{Nm(1)} \tag{39}$$

satisfies Equation (35). This threshold lies midway between the two conditional means of Z when the length N of the II buffer goes to infinity for fixed M_p; but when $N = M_p$, the threshold lies farther to the right by an amount proportional to the life extension factor. Equation (39) can be rewritten as

$$z'' = m(0)(1 + LM_p/NR) + m(1)/2 \tag{40}$$

where $R = m(1)/m(0)$ is the usual SNR.

CHAPTER 9

ENERGY EFFICIENCY IN STATISTICAL TESTING

When a communications signal is interjected on top of a stationary noise background, the receiver can be designed to come to attention when the short term average Equation (36) of the squared input rises above a threshold. By expending its limited energy during periods of alertness following this call to attention, the receiver can extend its service life beyond what would result in the absence of an alertness strategy. If the threshold is given by Equation (40) it will provide optimal performance subject to a prescribed life extension factor and some other parameters of interest. Yet the most critical parameters, which are those describing the relative rates of energy consumption in the quiescent versus the alert state, will be determined by a complex of engineering design decisions that are strongly influenced by circuit theory and device technology concerns that are not amenable to broad generalization.

A fundamental insight which is commonly employed in the design of isolated data processing systems is that CMOS integrated digital circuits draw power proportional to the rate of gate-level activity (as opposed to NMOS and bipolar devices in which quiescent power is not much less than the maximum). In a CMOS data processor, the arithmetic operations rate (AOR) will determine power consumption in a quasilinear proportionality. For very high AORs, there will be no great advantage over the other device types. The use of CMOS for energy conservation must be predicted on an algorithmic structure that minimizes the AOR in light of the computational requirement.

The term efficiency is used in the literature of mathematical statistics to refer to the relative data volumes required by two statistical tests to make the same decision at the same level of accuracy. There are various definitions of efficiency corresponding to different measures of accuracy. Casually speaking,

if for a given SNR the first algorithm (A) applied to N(A) data points gives the desired accuracy, and the second algorithm (B) would have to work on N(B) data points to get the same accuracy, then the relative efficiency of A to B is N(B)/N(A); and the higher this efficiency ratio, the better A is in comparison to B. In the design of algorithms for alertness in isolated systems, efficiency will be of some importance. Even if there is no shortage of data, the amount of time and energy required by any algorithm to operate on N data points will tend to increase as some power of N.

Perhaps the more critical question pertains to the nature of the arithmetic that the processor performs. Indeed, the time required to execute an FFT or some other fundamental transformation is usually taken as proportional to the number of multiply-and-add operations involved. Moreover, it is the multiplications that dominate the computation, since the digital multiply is realized as a sequence of sums. Thus an algorithm that consists in finding N(B) simple sums might be executed in a CMOS device for much less energy than an algorithm requiring N(A) simple products, even if N(A) is much less than N(B).

In order to compute the short term average of the squared input, a processor forms the sum of N products as stated in Equation (36). (Division by N is introduced for the convenience of the human analyst.) If t is the time (or energy) required to multiply, and t' is the time (or energy) required to add, the processor expends

$$C(1) = Nt + (N - 1)t'$$

units in the formation of the Z-statistic. Now suppose that the input (X) is provided through a J-level quantizer and that N is much greater than J. The N data points define the sample distribution

$$f(x,N) = n(x)/N$$

where $n(x)$ is the number of times x occurs in the sample. The test statistic can be reformulated as

$$Z = \sum_{x} x^2 f(x,N) \tag{41}$$

which needs only $Jt + (J - 1)t'$ units of time or energy. Thus the formulation of the Z-statistic based on the sample distribution is more "productive" by about a factor of N/J on the assumption that it takes negligible time or energy to sort X's into bins and compile the distribution. As a worst case, it might take Nt' units to define $f(x,N)$. If $t/t' = v$, the productivity of the formulation Equation (41) relative to Equation (36) is approximately

$$\frac{v + 1}{vJ/N + 1}$$

which approaches $N/J >> 0$ for large v.

Equation (41) defines the second moment of the distribution of the past N observations and thereby provides a statistic that can be used to discern the presence of a signal in the noisy channel. Other functions of the sample distribution can serve the same purpose. Given a priori knowledge of the distribution $p(x)$ conditioned on the null hypothesis, any of the measures of stochastic divergence listed in Chapter 6 might be employed. Indeed, the Z-test merely looks for a rise in the second moment of the sample distribution above a threshold set with reference to the second moment of $p(x)$. Information is lost in restricting attention to the second moment when the whole distribution is known (except when the distribution is Gaussian). Based on the results derived in Appendix A, the Kullback directed divergence, also known as cross entropy or relative entropy, is <u>more efficient</u> than the Z-test in a broad class of situations. Its productivity would appear to be about one-third that of the Z-test, if the time (energy) required to look up $\log f(x,N)$ in a table of logarithms is about the same as t'.

Maximum productivity might be attained in some processors by using the Kolmogorov-Smirnov statistic to measure stochastic divergence. This algorithm can be formulated non-parametrically, as a test of the stochastic equality of the last N observations to another sample of N observations made a various times across a broader epoch. It is shown in various statistical texts that the Kolmogorov-Smirnov statistic assumes the form of a sum of $N^2/2$ binary digits each of which indicates whether a given data point out of the last N is greater or less than a given data point from the broader epoch. A clever programmer could probably render this algorithm in a form that necessitates fewer than $N^2/256$ 8-bit sums. Regarding the efficiency of this test, if one assumes its power is the same as that of the Kolmogorov test from which it derives, the efficiency relative to the cross entropy test is on the order of $1/2\pi$ (typically). A detailed mathematical analysis might show that Kolmogorov-Smirnov test has better efficiency than the Z-test and better productivity than the cross entropy test; and that it is an optimal choice in systems reliant on certain types of hardware for specified channels.

CHAPTER 10

ARTIFICIAL INTELLIGENCE ASPECTS OF RECEIVER DESIGN

This chapter serves as an interim summary of the ideas presented so far. Chapters 1 through 6 dealt with the design of an adaptive signal detector, i.e., a device that uses its own continually updated "understanding" of the interference environment to adjust the form of the signal detection algorithm in an effort to maintain optimality as the environment changes. Chapters 7 through 9 considered the problem of "signal detection" in the (colloquial) sense of discerning the mere presence or absence of the signal in the noisy channel; and the significance of this problem in the context of isolated receiver systems was discussed. These two problems naturally complement each other. In so far as the whole discussion thus far has presumed that the receiver has access to the noise alone, unmixed with the signal, for the purpose of characterizing the interference environment, the tacit assumption has been that the channel is normally quiet (except for the noise) and that signalling occurs on a sporadic basis. Moreover, for the adaptive receiver to complete message reception prior to the next evolution in the state of the interference source, one must presume that message duration is generally short compared to the mean holding time of the interference source. If signaling is sporadic, then it makes sense to adopt an alertness strategy. Otherwise the bulk of the receiver's work consists in the vain exercise of trying to decode messages when only noise is present.

Figure 8 depicts a receiver structure of a generic type consistent with the suggestions described herein. The signal source sends A, B, C, or no signal; and its emission adds to the noise which is one of three types depending on the identity of the subsource operating at a given time. The sum sequence is stored in a receiver buffer of some fixed length. A statistical demultiplexer, defined in the same way as the SSD of Figure 4, controls an alert switch, turning on the main subsystem only when a test of stochastic equality denies the equivalence of

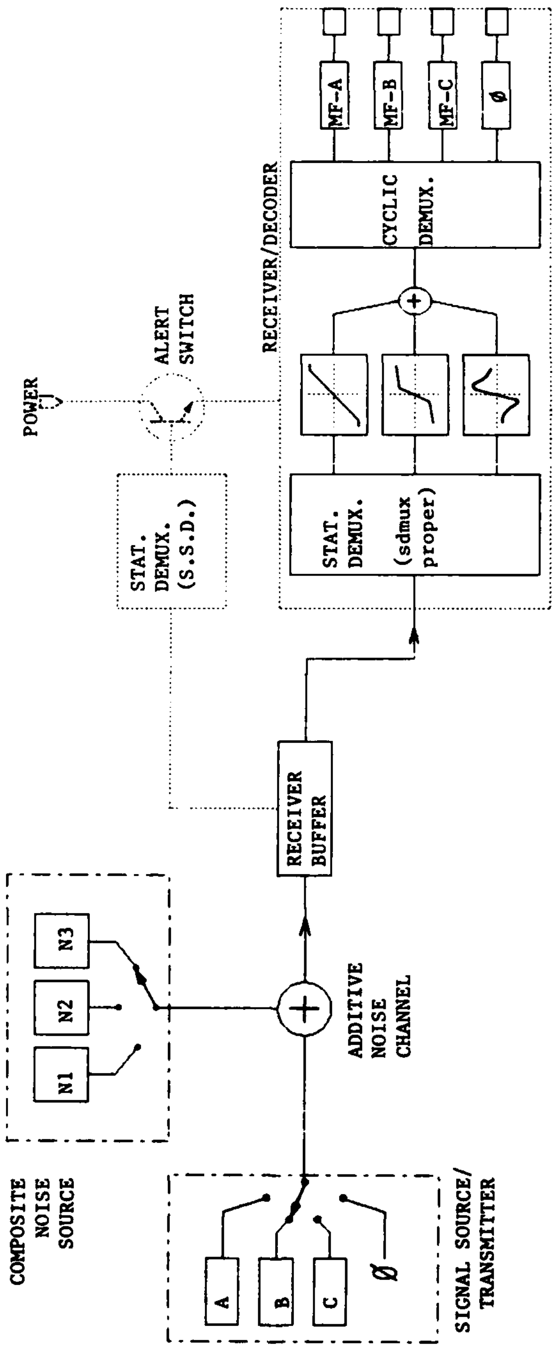

FIGURE 8. GENERIC COMMUNICATIONS RECEIVER WITH STATISTICAL DEMULTIPLEXING

the data in the buffer to the recent history of the input. A statistical demultiplexer (proper), understood in the context of Chapter 2, now classifies the state of the prevailing interference and decides which likelihood ratio test is the appropriate one. The liklihood ratio test for detection of the weak signal is equivalent to passing the input through a no-memory nonlinearity followed by a matched filter, as noted in Chapter 1. A cyclic demultiplexer routes the data from the correct nonlinear section through the bank of matched filters in succession. The diagram is somewhat ambiguous in that it does not explicity show that the sdmux (proper) has access to the longer-term sample which serves to define the prevailing noise prior to injection of the signal.

The two properties that Figure 8 is intended to embody are alertness and adaptivity. Appealing to the layman's understanding of psychology, and seeking analogies between the receiver's performance and mental processes, the antonyms for these two terms might be paranoia and perseveration, respectively. Paranoia is a term applied to individuals or groups afflicted with uncontrolled fears or suspicions which are not grounded in objective reality. The paranoid individual may be propelled into states of alert tension by the desire to clarify perceived threats that exist only in the imagination. In a similar manner, the isolated receiver that lacks an alertness strategy squanders its energy on the search for messages in the random processes of the environment that affect its sensors. The healthy human individual behaves under the control of a reticular activating system, physically centered in the brain stem, which is responsible for the mental phenomenon of selective attention. The reticular activating system is widely interconnected with the various parts of the brain; and it figures in the abilities to concentrate and to sleep.[12] Certain aspects of the ability to learn and adapt are associated with parts of the cerebral cortex located in the frontal lobes. Damage to the cortex in these areas can result in perseveration, a type of behavior that carries perseverence to the point of absurdity [Ibid]. Perseveration essentially consists in the "refusal" to learn the new rules of a game when the evidence overwhelmingly demonstrates that the rules have changed. In the Darnafalski game of Chapter 5, if there were two pairs of dice, one fair and the other weighted so that the number 12 showed up over half the time, a player who could not learn to guess which pair of dice is in use over the course

of many rounds would typify perseveration. Likewise, an engineer who designs an ELF receiver for optimal performance against a type of noise recorded on one occasion, when faced with the fact that his system performs poorly in a variety of situations, would deny the loss of optimality only by rejecting the overwhelming evidence.

CHAPTER 11

ADAPTIVE RECEPTION WITHOUT MULTIPLICATION

Some degree of alertness can be attained by a system in which the full-time processor (i.e., the SSD) performs no multiplications. This was pointed out in connection with the Kolmogorov-Smirnov statistic. A random sample of N observations drawn from the preceding B (B >> N) yields a p(x) representing the first order density of the noise at the present time on the assumption that a state transition has not occurred in the last B. The last N samples, however, yielded f(x,N). The Kolmogorov distance D from p to f will have a specified distribution under the null hypothesis just assumed. But if the signal appears in the last N, D will be too large to typify the null hypothesis; and the receiver will go to the alert state to extract the indicated information.

The optimal receiver can now employ p(x) to shape the nonlinearity through which the N most recent observations are sent in succession to the matched filter, which performs N inner product steps in the computation of a test statistic which is specific to one particular signal in the lexicon of the transmitter (which is known beforehand). In the notation of Chapter 9, the time required for each inner product step is t + t'. The time required to test lexicon for each of M distinct signals is MN(t + t') or greater. Consider the following question: Is there a way to accomplish this M-ary hypothesis test without doing any multiplications? If so, the algorithm would need about MLt' time units to execute. If L is much less than Nt/t' = Na, then it will be faster.

Define M sample distributions based on the last N observations by

$$f(x,N;m) = (1/N) \sum_{n=1}^{N} I\{x-[X(n) - s(n,m)]\}$$

where I is the indicator function which assumes unit value only when its argument vanishes and $s(n,m)$ is the n-th component of the m-th signal vector. There are N^2 simple sums involved in the computation of each of these M functions. For every m, compute the Kolmogorov distance

$$D(m) = \sup_x \left| \sum^{x} [f(x', N;m) - p(x')] \right|$$

from the p-distribution. The distribution of $D(m)$ appears to be the same for every m so long as $s(n,m)$ is an uncorrelated sequence and

$$S = \sum_n s^2(n,m)/N$$

is the same for every m. In fact, in the limit as N approaches infinity,

$$\lim D(m) = \sup_x \left| (1/2) Sp'(x) \right| , \quad m > 0 , \tag{42}$$

under the null hypothesis, where p' is the derivative of p with respect to x. For example, if p is a Gaussian density, with variance v, then the limit reduces to

$$\frac{\text{"SNR"}}{8.264}$$

where SNR = S/v. In the same limit, the mean of D(0) under the null hypothesis goes to zero as the reciprocal square root of N. Here the case m = 0 refers to the signal consisting of all zeros. It is assumed without proof that D is also root-N consistent for the M proper values of m, converging to the indicated limit under the null hypothesis.

If the signal indexed by m' is present, then D(m') will go to zero as D(0) did under the null hypothesis (H0: m = 0), since D(m') is generated by a sequence of random variables having the same character as the input under H0. In this case, D(0) will have the same limiting expected value as given by Equation (42). For the other cases

$$\lim D(m) = \quad 2 \lim D(m'), \; m = m' \text{ or } 0 \; .$$

Thus "the hard part" is to discriminate against H0 in favor of Hm' when m' is true.

The optimal procedure dictated by the Neyman-Pearson lemma would generate M + 1 test statistics each of which is asymptotically normal under the null hypothesis and each of the M alternatives. The variances of these normal statistics are proportional to 1/N and the spreads between the means converge asymptotically to values that contain the first power of the SNR as a common scale factor. Comparing these to the D-statistics just described leads one to conclude that the asymptotic efficiency of the latter relative to the optimal does not vanish. Resort to reception without multiplication therefore involves a controlled (as opposed to catastrophic) loss of fidelity that can in general be compensated for by reducing the (design) data rate.

The Kolmogorov distance is only one of the measures of stochastic divergence than could be applied to the residual sequences producted by subtracting the signal in question (point-by-point) from the input. The Kullback divergence would engender a receiver algorithm which us superior within this class of algorithms that try to find stochastic equivalences involving residuals.

REFERENCES

1. Antonov, O., "Optimum Detection of Signals in Non-Gaussian Noise," Radio Engineering and Electronic Physics, Vol. 12, 1966.

2. Evans, J. E., and Griffiths, A. S., "Design of a Sanguine Noise Processor Based on Worldwide ELF Recordings," IEEE Trans. Communications, Vol. COM-22, 1974.

3. Baran, R. H., Adaptive Signal Detection for the Optimal Communications Receiver, NSWC TR 83-236, 1983.

4. Kassam, S. A., "A Bibliography on Nonparametric Detection," IEE Trans. Information Theory, Vol. IT-26, No. 5, 1980.

5. Gallager, Information Theory and Reliable Communication, (New York, NY: Wiley, 1968), circa page 67.

6. Berger, T., Rate Distortion Theory: A Mathematical Basis for Data Compression, (Englewood Cliffs, NJ: Prentice-Hall, 1971).

7. Bickel, P.J., and Doksum, K.A., Mathematical Statistics, (San Francisco: Holden-Day, 1977).

8. Cheng, P.E., and Serfling, R.J., "Asymptotic Mean Integrated Squared Errors of Some Nonparametric Density Estimators," IEEE. Trans. Information Theory, Vol. IT-27, No. 2, 1981.

REFERENCES (Cont.)

9. Toussaint, G.T., "Sharper Lower Bounds for Discrimination Information in Terms of Variation," IEEE Trans. Information Theory, Vol. IT-21, No.1, 1975.

10. Brockett, P.L., et. al., "Information Theoretic Analysis of Questionnaire Data," IEEE Trans. Information Theory, Vol. IT-27, No. 4, 1981.

11. Birnbaum, Z. W., "Numerical Tabulation of the Distribution of Kolmogorov's Statistic," J. Amer. Statistical Association, Vol. 24, p. 467, 1953.

12. Calvin and Ojemann, Inside the Brain, Times Mirror, New York, 1980.

APPENDIX A

THE POWER OF THE CROSS-ENTROPY TEST

Let the input be

$$Y(n) = X(n) + C[n] \tag{A-1}$$

with $\{X(n), n = 1, 2, \ldots\}$ and iid sequence and $(C[1], \ldots, C[N])$ belonging to a set A of M distinct N-vectors or consisting of N zeros. The signal detector (SD) is defined as any procedure which discriminates between the null hypothesis

$$H(0): \quad C[n] = 0 \text{ for every } n$$

and the composite alternative

$$H(+): \quad C \in A .$$

Assume that either H(0) or H(+) must hold for $n = 1, 2, \ldots, N$; and that the elements of A are constrained by equalities of the form

$$(1/N) \sum_{n=1}^{N} (C[n])^{t} = w(t) \tag{A-2}$$

for known moments $w(t)$. The signal classifier (SC) is defined as a procedure that classifies the input $\{Y(1), \ldots, Y(N)\}$ as representing the m-th element of A or else asserts H(0). In other words, the SC performs an M-ary hypothesis test to select one of the following:

$$H(0): \quad \underline{C} = \underline{0}$$

H(1): $\underline{C} = \underline{C}(1)$

H(2): $\underline{C} = \underline{C}(2)$

. . .

H(M): $\underline{C} = \underline{C}(M)$,

with the presumptions that $\underline{C}(1)$ through $\underline{C}(M)$ are nonzero and

$$\underline{C}(m) - \underline{C}(m') = 0$$

if and only if m = m'. Now if the SD controls the SC to the extent that the SC operates on the input only when SD asserts H(0), to what extent can SD be simpler, faster, or less computation-intensive than SC without incurring an unacceptable probability of false rest?

Let the SC use the optimal procedure for classifying the input. Defining $\underline{C}(0) = \underline{0}$, the output of the SC is m, the best estimate of m ϵ (0, 1, ..., M). The maximum likelihood estimate is obtained by selecting the supremum of the M + 1 likelihood functionals

$$L(m) = \sum_{n=1}^{N} \log\{p[X(n) - C(m,n)]/p[X(n)]\} \tag{A-3}$$

When the average C^2 (m,n) is much less than EX^2, the approximate form of Equation (A-3) is

$$L(m) = \sum_{n=1}^{N} C(m,n)[-\log p(x)]'_{x = X(n)} \tag{A-4}$$

where the prime indicates the derivative of log(p) with respect to x. If

$$-d[\log p(x)]/dx = g(x) \tag{A-5}$$

is a known function, then the computation of $L(0)$ through $L(M)$ appears to require NM floating-point operations (flops) per input. In other words, if the input is obtained by sampling a continuous data stream at a rate of R times per second, the SC has to be able to do NMR flops/second in order to operate in real time.

The SD must be able to indicate the presence of signals in the additive noise channel, with suitable detection and false alarm probabilities, while performing much fewer than NMR flops per second. Consider a class of SD's whose work rate is vNR flops/second, where $1 < v < M$. When $w(0) = 0$, the presence of the signal is manifested by increased average power at the output of the channel:

$$E[Y^2 | H(+)] = E[Y^2 | H(0)] + w(2) , \tag{A-6}$$

where E takes an average over all the $Y(n)$. In fact, the sum of the last N values of Y^2 is asymptotically normal with its mean shifted $Nw(2)$ units to the right by $H(+)$. Moreover, this procedure would seem to be most powerful for testing $H(0)$ against $H(+)$. The work rate is clearly on the order of NR (so that $v = 1$).

Now another class of SD's will attain work rates of the form vNR, with $0 < v < 1$, by classifying the input as one K discrete values, where $K << N$, and by performing a test based on the distribution of these quantized samples. Information is lost in quantization and the power of the test procedure will suffer to some extent. Let the output of the quantizer be $Z(n) \epsilon \{z\}$, the cardinality of this set being $|\{z\}| = K$. Then

$$E\{[Z(n) - Y(n)]^2 | H(m)\} = M.S.E. \tag{A-7}$$

is the mean squared error generated in quantization. The mean square error is minimized, subject to some assumptions about the source statistics, using rules developed by Lloyd[A-1] and others. In particular, the quantization may be finer over the range of more probable input values, and coarser over the range of less probable values. Lloyd shows conditions in which the width of an

interval should be proportional to the cube root of the probability density of the signal (or ensemble of signals). Whether optimal or not, one expects the distribution of

$$T = [Z^2(1) + \dots + Z^2(N)]/N - E[Y^2 | H(0)] \tag{A-8}$$

to be asymptotically normal with conditional means of w(2) and zero (under H(w) and H(0), respectively) and a common variance given by

$$\mathrm{Var}(T) = \mathrm{M.S.E.}^2/N + \mathrm{Var}[Y^2 | H(0)]^2 = \mathrm{M.S.E.}^2/N + 2(EX^2)^2/N \tag{A-9}$$

when the M.S.E. is independent of m and w(1) = 0. (I.e., it is assumed that C(n,m) has the same symmetric distribution for every m, and that X is likewise symmetrically distributed about zero.) If T(a) is the threshold that gives an a-level test, so that

$$T(a): \ \ \mathrm{erfc}\{T(a)/[\mathrm{Var}(t)]^{1/2}\} = a \ , \tag{A-10}$$

then the power of this test is simply

$$P[a] = \mathrm{erfc}\{[T(a) - w(2)]/[\mathrm{Var}(T)]^{1/2}\} \ . \tag{A-11}$$

But W is computed by the rule

$$NT = n(1)z^2(1) + \dots + n(K)z^2(K) \ , \tag{A-12}$$

where n(k) inputs were quantized as z(k),

$$n(1) + \dots + n(K) = N \ , \tag{A-13}$$

and the z(k) are the ordered elements of {z}. Therefore, if the $z^2(k)$ are fixed (stored) numbers or if they vary slowly enough to require re-computation only rarely, the SD needs to do only K multiplications, and the work rate is vNR when v = K/N << 1.

Another way to achieve $v = K/N$ would be to compare the non-negative cross entropy statistic to a threshold and reject H(0) when the threshold is exceeded. The cross entropy of a discrete density q(k) with respect to a discrete density p(k) on the same range {z} is

$$I(\{q\};\{p\}) = \sum_{k=1}^{K} q(k)\log[q(k)/p(k)] \ , \tag{A-14}$$

where $q(k) > 0$ for every k and $0\log0 = 0$. With reference to (A-13), take $n(k)/N = q(k)$, although the real time program would not bother to normalize using fixed N. The density p(k) in Equation (A-14) describes the distribution of Z(n) under H(0). When H(0) is true, the distribution of NI (being N times the cross entropy statistic) is asymptotically chi-square with K/2+1 degrees of freedom according to Brockett[A-2] who cites proof by Kullback. When H(+) holds, take

$$f(x|+) = \sum_{(z)} r(z)f(x - z) \tag{A-15}$$

for the (continuous) density of X, where

$$r(z) = \sum I^{+}(\{C(n,m) = z\}/N$$

and I^{+} is the indicator function. Expanding Equation (A-15) in a Taylor series to second order,

$$f(x|n) \doteq f(x) + (1/2)w(2)f''(x) \tag{A-16}$$

where f(x) is conditioned on H(0) and, again, $w(1) = 0$ in Equation (A-2). Now the expression $\log[q(k)/p(k)]$ in Equation (A-14) corresponds to

$$\log[f(x|+)/f(x) = \log\{[f(x +) - f(x)]/f(x) + 1\} \tag{A-17}$$

$$\approx f(x|+)/f(x) - 1$$

Substituting Equation (A-16) in Equation (A-17), multiplying by Equation (A-16), and integrating, one has

$$E[I|H(+)] = (1/4)w^2(2)\int [f(x)]^{-1}[f''(x)]^2dx \qquad \text{(A-18)}$$

for the mean of the cross entropy statistic conditioned on H(+), in the limit of vanishing w(2). A similar argument leads to the conclusion that the asymptotic distribution of NI under H(+) is chi-square with noncentrality parameter given by N times Equation (A-18). When K is large, the chi-square random variable with K/2+1 degrees is approximately normal with mean of K/2 and variance of K. Hence, the definition

$$I(a): \quad F(NI(a);K/2) = 1 - a , \qquad \text{(A-19)}$$

with F(x;K/2) the chi-square df for K/2 degrees, leads to the result

$$P_I(a) = \operatorname{erfc}(N[I(a) - hw^2(2)/4]/K^{1/2}) \qquad \text{(A-20)}$$

for the power of the a-level cross-entropy test where, in accordance with Equation (A-18), h is determined by the distribution of X.

Referring back to Equation (A-11), the power P[a] of the a-level second moment test was given by that expression which is essentially the same as

$$P[a] = \operatorname{erfc}\{N^{1/2}[T(a) - w(2)]/EX^2\sqrt{2}\} . \qquad \text{(A-21)}$$

Comparison to Equation (A-20) shows that $P_I[a]$ may exceed P[a] when the coefficient of variation

$$R[I] = Nhw^2(2)/K^{1/2}$$

is significantly greater than

$$R[T] = N^{1/2}w(2)/EX^2\sqrt{2} .$$

Then for

$$(N/K)^{1/2} hw(2)EX^2\sqrt{2} >> 1 ,$$

the I-test is better; and since

$$N/K = 1/v >> 1$$

by design,

$$\sqrt{2}\, hw(2)EX^2 >> v \tag{A-22}$$

is the condition to be satisfied. In general, h will be of the form

$$h \approx b/(EX^2)^2 , \tag{A-23}$$

where the dimensionless factor b depends more on the shape or functional form of the distribution of X than on the scale factor that figures in the variance, EX^2. Then Equation (A-21) reduces to

$$\sqrt{2}\, b >> v/S.N.R. , \tag{A-24}$$

where S.N.R. = $w(2)/EX^2$ is the usual signal-to-noise ratio. The R.H.S. of the last inequality is the ratio of two dimensionless numbers which are individually less then unity.

REFERENCES

A-1. Lloyd, S. P., "Least Squares Quantization in PCM," IEEE Trans. Info. Theory, Vol. IT-28, No. 2, Mar 1982.

A-2. Brockett, P. L., et. al., "Information Theoretic Analysis of Questionnaire Data," IEEE Trans. Information Theory, Vol. IT-27, No. 4, 1981.

DISTRIBUTION

	Copies
Office of Naval Research Attn: Technical Library Arlington, VA 22217	2
Commanding Officer Office of Naval Research Branch Office 495 Summer Street Boston, MA 02210	1
Commanding Officer Officer of Naval Research Branch Office 536 South Clark Street Chicago, IL 60605	1
Commander Naval Air Systems Command Attn: Technical Library Department of the Navy Washington, DC 20361	1
Commander Naval Sea Systems Command Attn: Technical Library (99612) Washington, DC 20362	2
Commander Naval Electronics Systems Command Attn: Technical Library Department of the Navy Washington, DC 20360	1
Commander Naval Ocean Systems Center Attn: Technical Library San Diego, CA 92152	1
Library of Congress Attn: Gift and Exchange Division Washington, DC 20540	4
Commanding Officer Naval Underwater Systems Center Attn: Technical Library Newport, RI 02844	1
Commander Naval Weapons Center Attn: Technical Library China Lake, CA 93555	1
Superintendent Naval Postgraduate School Attn: (Curricular Officer for Weapons, Engineering, and ASW) Monterey, CA 93940	1
Commanding Officer Naval Research Laboratory Attn: Code 1000	1
Technical Library Washington, DC 20375	1
Defense Technical Information Center Cameron Station, Building 5 Alexandria, VA 22314	12
Internal Distribution:	
U12 (R. H. Baran)	10
U12 (J. D. Sherman)	1
U14 (S. Keadle)	1
U21 (P. Jackins)	1
U21 (S. Hebbert)	1
R04 (D. Love)	1
R43 (J. Coughlin)	2
R42 (W. Caswell)	1
K34 (T. Orlow)	1
N21 (M. Masters)	1
G52 (Sholander)	1
G23 (Ohlmeyer)	1
E231	9
E232	3

END

FILMED

7-85

DTIC